Lauren Huret

ARTIFICIAL FEAR INTELLIGENCE OF DEATH

In conversation with Monica Anderson,
Erik Davis, R.U. Sirius and Dag Spicer

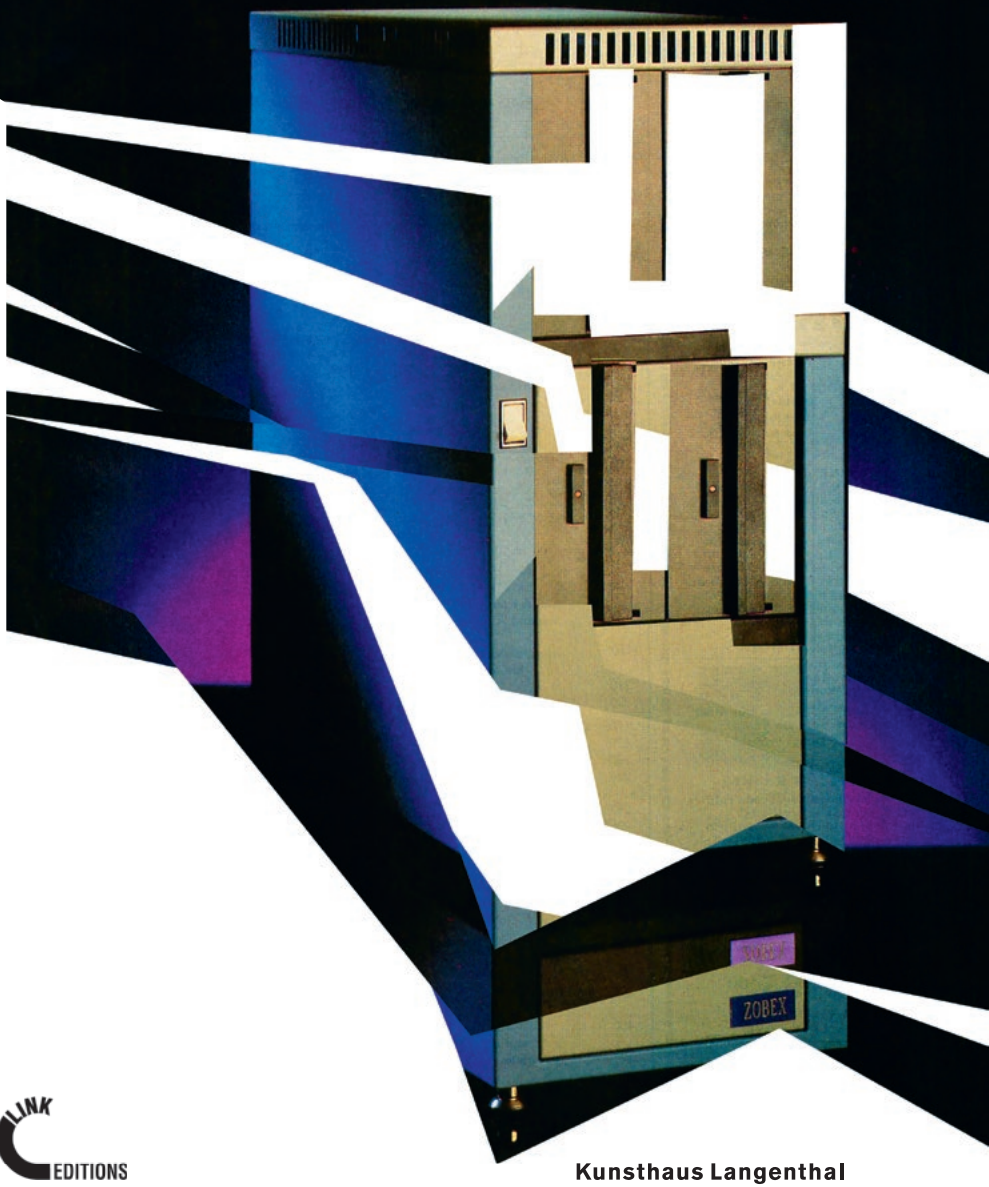

LINK EDITIONS

Kunsthaus Langenthal

Lauren Huret

ARTIFICIAL FEAR INTELLIGENCE OF DEATH

In conversation with Monica Anderson,
Erik Davis, R.U. Sirius and Dag Spicer

LINK EDITIONS

Kunsthaus Langenthal

Lauren Huret
Artificial Fear Intelligence of Death
In conversation with
Monica Anderson, Erik Davis,
R.U. Sirius, Dag Spicer

Published by LINK Editions,
Brescia, 2016
www.linkartcenter.eu
Co-published by Kunsthaus
Langenthal
www.kunsthauslangenthal.ch
ISBN: 978-1-326-61268-9

Edited by: Kunsthaus Langenthal
Concept: Lauren Huret, Raffael Dörig,
Hunter Longe
Editorial office: Raffael Dörig,
Claire Hoffmann
Transcription and proofreading: Hunter
Longe, Erik Smith
Graphic design: P'INC. AG, Langenthal
Print of numbered special edition:
Merkur Druck AG, Langenthal
The print-on-demand version is printed
and distributed by www.lulu.com

Published on the occasion of
the exhibition of the same
title at Kunsthaus Langenthal,
April 27 – June 26, 2016
Curator: Raffael Dörig
Assistant curator: Claire Hoffmann
Exhibition secretary: Lotti Vonäsch
Technical assistance: Brigitte Jost,
Marc Masson, Erwin Aktasli

The artist would like to thank: Monica
Anderson, Erik Davis, Ken Goffman,
Dag Spicer, Hunter Longe, Raffael
Dörig, Claire Hoffmann, Nicole and
Sylvain Huret, Nan Hill & Brian Longe,
Paul Bernard, Megan Prelinger,
Celeste Baranski (Numenta) FCAC
de Genève and all my friends.

This publication was made possible
by the generous support of

STIFTUNG JABERG LANGENTHAL

prohelvetia

MIGROS
kulturprozent

AVEC LE SOUTIEN DE LA VILLE DE GENÈVE

REPUBLIQUE ET CANTON DE GENEVE

Avec le soutien du Fonds cantonal
d'art contemporain, SCC-DIP, Genève

Partner Kunsthaus Langenthal:
Ammann Unternehmungen,
Langenthal; création baumann,
Langenthal; MOTOREX-
Schmiertechnik, Peter R. Geiser,
Langenthal, GÜDEL AG, Langenthal;
Merkur Druck AG, Langenthal;
P'INC. AG, Langenthal.

The Kunsthaus Langenthal is sup-
ported by the Canton Bern, the City
of Langenthal and the communities
of the Regional Culture Conference.

The many examples of applied artificial intelligence, or "weak AI," that now populate our lives (Siri, Internet search engines, chatbots, auto-translate, facial recognition, etc.), developed by major tech companies and secular institutions, foreshadow the presupposed future appearance of a form of non-organic autonomous intelligence. Such abundant use of weak AI is very recent but sociological and behavioral studies suggest we have a "natural" empathy and powerful emotional reaction vis-à-vis these objects/programs and their representations (user interfaces).

In fact, forms of non-organic intelligence are already surrounded by stories, behaviors, and personifications as well as political and social claims. We must constantly face, integrate, and learn what functions they serve and the phenomena associated with them. We must continually try to make sense out of what we may feel, live, and experience with, in, and around these machines with the help of narratives, mythologies, re-appropriations, and interpretations, like the ones manifested in the latest science fiction movies. AI has an undeniable attraction on our imagination, and it opens an infinite field of exegesis, interpretative shifts, indescribable confusion and stupor. Despite the many philosophical questions that arise (e.g. What is intelligence?), tremendous breakthroughs in the last few years have brought AIs much closer to us.

Beginning in 2013, I started delving into the myriad of technical, philosophical, and sci-fi literature surrounding AI as well as its role on the internet. In the summer of 2015, with support from the FCAC Geneva, I undertook a trip to the San Francisco Bay Area and the epicenter of computing: Silicon Valley, California. I wanted to talk face to face with thinkers who could offer varying perspectives, profound criticism, funny, scary, and humbling insights, and an overview of artificial intelligence. For this publication, I have selected

the interviews that I conducted with Monica Anderson, Erik Davis, R.U. Sirius, and Dag Spicer because they reflect a multiplicity of viewpoints and the complexity of AI. During my visit to the Prelinger Archives in San Francisco, Megan Prelinger introduced me to *Byte* magazine, an important and popular monthly magazine from the early days of personal computers. Visual material from issues of *Byte* from 1978 to 1986 forms the basis of the collages in this book. When looking through this iconic material, it was evident that the invention of personal computers was about to impact every level of society in a major way. Often compared to a magical object with fantastic abilities, the computer has been seen—since its inception—as a "miracle." From the same source I chose a selection of advertisements in which computers are ascribed human or magical qualities, resonating, often in a goofy and playful way, with the notion of a human-like computer that is so central to AI discourse.

I invited Hunter Longe to collaborate and assist in my knowledge-gathering excursion, of which this book is one of the outputs. His voice is present in several of the interviews. Without him the project would not have been possible, not only because he is originally from the Bay Area but because his knowledge on the topic made him an irreplaceable dialog partner. I would like to thank him, his family, and his friends who have been so kind and supportive. I would also like to thank the interviewees who were extremely generous with their time and thoughts, as well as Raffael Dörig and Claire Hoffmann for all of their work and for believing in the project.

Lauren Huret, April 2016

INTERVIEW
DAG SPICER
I DON'T NEED SCIENCE FICTION, BECAUSE I AM LIVING IT EVERY DAY

Lauren Huret: *Dag Spicer, you're the Senior Curator of the Computer History Museum in Mountain View, in the heart of Silicon Valley. You have a background in the history of technology and the history of science, but you started out as an electrical engineer, as a digital circuit designer. I imagine this has shaped your view on the history of computing as well?*

Dag Spicer: A digital circuit designer is someone who works with transistors and integrated circuits and resistors and diodes and capacitors—all of these electronic components—and puts them together in certain ways to build systems. I do it as a hobby. When I was really engineering, the packaging for integrated circuits was known largely as "through-hole" packaging, meaning that the pins of the IC went through the circuit board and were soldered underneath the board. About twenty-five years ago, a new technology called surface-mount technology arrived and that does not have any holes, the chip is laid on top of the board and the wires run on the surface. You can populate the board on both sides so you can double the density of the circuit board. Because the surface-mount chips are so much smaller than a through-hole chip, you can put more of them in there—surface-mount technology is what actually enabled things like mobile phones. So, the ability to shrink the size of the circuit is not only an effect of Moore's law, it's also related to other factors like packaging. As an example, there is a new chip that just came out. It's a voltage regulator and its size is one square millimeter! You are supposed to solder that, but you can barely even see it. So the obvious factor about surface-mount technology is that it's really made for robots, not for human assembly. That's a huge change as well. Once you get robots building computing devices, you can make them smaller and have better control over the pitch and the fineness of lines and how close components can get together and so on. You increase the density.

LH: *So building circuits is no longer for humans?*

DS: You can prototype your circuit on a bread board, a kind of electronic *tabula rasa* on which you can interconnect the components of your circuit quickly just to see if it will work. But once you get into production, you advance to surface-mount construction and that obviously requires other skills. Probably a circuit board designer for example, manufacturing engineers, quality control people. Dozens of different skill sets are required once you get into production.

LH: *Can you tell me more about computer architecture?*

DS: Computer architecture is the study, at a high level of abstraction, of how computers are built. The idea of an architecture is to be something that lays out the basic blocks of a computer system. Computer architects can mix and match these blocks to give the computer certain performance characteristics. To take a historical example, the computing market used to be segmented into scientific and business computers. If you were designing a scientific computer, it would usually do elaborate floating-point calculations and so the architect would optimize the architecture of the computer for floating-point computation by mixing the major blocks in a certain way. If your computer was designed for a business market, customers deal with numbers that are neither extremely small nor extremely large, as with a scientific computer. Often, these users require only a simple architecture and so there are fewer (and slightly different) blocks in that architecture. You can think of computer architecture like a blueprint for a building. Architecture is actually independent of materials. You can build a building out of brick, out of wood, out of stone, out of plastic, out of metal, but the plans could look the same. For example, at the Computer History Museum, we have several pieces of the massive 1946 ENIAC computer, a computer so large that you can walk inside it.

In 1995, students at the University of Pennsylvania re-implemented ENIAC using a VLSI integrated circuit as a class project. The ENIAC architecture remained the same; only its implementation (the materials used to make it) changed: the ENIAC used vacuum tubes and the chip uses transistors but it's the same architecture exactly.

Probably the best computer architecture ever designed was for the IBM System/360, which came out in 1964. It was a family of computers spanning a fifty-to-one performance range. If you were a small company you could lease or buy the small model; if you were a big company you chose one of the bigger models. All models ran the same software, which is really great because as a business grew it could get a more powerful model of the System/360, yet still preserve its software investment. Nobody did that before IBM. Usually companies had to throw out much of their software and start all over again every time they changed computers, especially if they changed the manufacturer, as was the case if you switched from IBM to Honeywell or from Honeywell to RCA or General Electric. So the System/360 architecture was designed to offer software compatibility across a very wide range of performance, and these performance goals were achieved by implementing the same architecture in different technologies and using different "building blocks" as I noted above. The current series of IBM mainframe computer, called the Z-series, still runs System/360 software—more than half a century later.

That's really important because IBM mainframes run some of the largest businesses in the world (like American Express, Visa, banks, government, industrial concerns etc.). They don't want to change their software. There is too much money at stake. Every time they make changes it could ruin them if they are not careful. So they don't want changes to their software, they would rather keep running the same old software but on newer and newer hardware and that's what they do. You can actually run an IBM

System/360 environment on your laptop if you use an emulator-software that lets you fake one computer on another. There is a Unix-based emulator for the System/360 and it looks exactly like an IBM System/360 mainframe from 1964. You can run all the same software. In the history of computing community, we rely very heavily on this technique of emulation because it's so hard to keep the original hardware operating. It requires a lot of tacit knowledge, knowledge that is not written down—it's often in the heads of the engineers that fix it, they'd say, "Oh yeah you got to kick it over there for it to start working." I call it a recipe. The recipe for keeping these things going has been forgotten or lost or people have retired. So it only makes sense to use current platforms: Unix, Windows, or Mac, and emulators. There are dozens and dozens of historical machines, most of which we have, that are huge, they are the size of a refrigerator or even a room, and you can emulate them on a PC just through some software. So that helps us preserve it—not so much the user experience, but at least, you can see how the software functioned. The emulators even create virtual magnetic tape drives in software that you can load programs onto, as well as virtual disk drives and printers that you can use just like in a real mainframe installation. So the computer in this sense is being used to preserve the history of computing, and emulation is one of the key methods.

LH: *In your bio there is a phrase stating that your motivation to study the history of technology and science is about, "seeking to understand the deeper forces at work behind the nature of scientific and technological innovation." So, what are those driving forces, especially here in Silicon Valley, and how did this area become the epicenter of computing?*

DS: It's hard to know how to break that down, because on the one hand it all goes back to Moore's law and the fact that we get more and more powerful platforms to play with.

Worrying about paper data is now an old story hidden under the table

Sunbath

Those platforms give us options software-wise. You can't write a GPS application when you have 4 KB of memory like in a computer from 1965, but when you have 4 GB of memory, all of a sudden the application space opens up. Today the web has become the single most important aspect of computing. Nobody saw that coming fifty years ago. If I look at Silicon Valley, clearly there are two or three key elements. One is the huge venture capital community here, which loves taking risks. They don't care about losing money it seems. They are willing to bet on ten companies and if one of them succeeds, that's considered okay. The other thing is excellent universities, mainly Stanford and Berkeley, which are world-class. And don't discount the weather. I'll give you a little story: William Shockley, the co-inventor of the transistor, won a Nobel prize in 1956 and shared it with his two co-workers, John Bardeen and Walter Brattain. He moved to Mountain View and set up his factory there because he wanted to be close to his mother who lived in Palo Alto. He was the single most important figure in semiconductor electronics and he brought the world's best chemists and physicists in the world to his little laboratory in Mountain View—people like Gordon Moore and Robert Noyce. I often wonder, if he hadn't come here would there even be a Silicon Valley? Because from that company Fairchild was spun off, and from Fairchild about fifty other small semiconductor companies were spun off, and from those companies you had still more spin-offs. So there are hundreds of companies that trace their origins back to the Shockley Semiconductor Laboratory. So you have the venture capitalists, the beautiful weather, the great universities and this one odd variable that the key figure in semiconductor electronics moved to the area and attracted the other top people in that field to this area. There are a lot of regions around the world that have tried to replicate what

that magic formula is. No one has been able to replicate Silicon Valley, but they have come up with pretty good ones. Ireland has very good industry and they have a favorable tax regime there. And Germany is always amazing.

LH: *As a historian, can you give any examples of early artificial intelligence projects?*

DS: The first thing was to develop a language for AI and that turned out to be LISP. It was developed by John McCarthy in the late 1950s at MIT. One of the early important applications written in LISP was Dendral, starting in 1965 and developed by Edward Feigenbaum and others. It was one of the first "expert systems," emulating the decision-making of a human expert. LISP is still used for AI. It's sixty years old and it's still the dominant and probably best language for AI.

LH: *There seems to be a real surge of artificial intelligence start-ups in the last few years, as well as Hollywood films and, I would say, an Internet paranoia. According to you, what is driving this recent re-interest in AI?*

DS: One of the things that's driving interest in AI is the wide dissemination of cheap platforms to develop AI applications on, whether they are little robots or just software environments. The major players like Google, Apple, and Facebook are providing development kits and development environments for their platforms like Android and Oculus. They give people and developers the tools they need to create applications for free, to sell more hardware. So it's the availability of cheap and easy development kits and APIs, the Application Programming Interfaces, which are the interfaces that companies make to their products. The open-source community is involved in a lot of that stuff. I wouldn't say Arduino is big in the AI community but it's a similar kind of thing. It's a very popular platform, it's cheap as dirt and there are tons of people developing on it.

I saw a chip—so this isn't software, it's actually a chip—that you design into a camera and the chip provides what's called emotion recognition. So it would look at a face and say you are surprised, you are sad, you are angry, you are happy. That's AI. It's using artificial intelligence algorithms on this tiny chip. I'm not sure what the application is. Maybe if you want to filter by, let's say, only people who are smiling, or sad … I'm really not sure what the application is, but that's so typical of engineers. They design things that have no application just because they think it's neat (laughs).

LH: *Do you think we will ever understand the brain?*

DS: My quibble with the AI community is that they are relying on the principle of "emergence" to explain how a machine might develop consciousness. Their basic idea is that if you build enough neurons and interconnections, at some point, self-awareness just happens. To me this begs the question what consciousness is and is no more sophisticated an argument than biblical beliefs about God breathing life into a lump of clay. According to many AI researchers, at some point consciousness just "happens." While this has worked at small scales in "swarm intelligence" models of behavior, it's a long leap from fifty robots self-organizing into a group to creating a human (or human-like) brain. It's this *deus ex machina* approach to AI that makes me skeptical.

LH: *What does intelligence mean to you and what do you think these days about the field of machine intelligence?*

DS: As far as the state of machine intelligence is concerned, it's very far advanced in certain domains. The general project of making a human-like brain is a non-starter I my opinion, it's a Holy Grail. I see it in a 500-year time span, maybe we'll get there. But companies have made great progress in niche applications (like Google Maps or Siri). It's like you

Man's intelligence trapped in computers

don't even think of it as AI anymore. You know you can talk to your iPhone, the voice-to-text dictation, it's sending your voice out over the web in real-time and translating it, that's AI. And it will also do it in French and it gets all the accents correctly; that is magic to me. And often, if you come to technology and you think, "Wow, that's really magic …" there is probably AI involved in it somewhere. To me, I guess an intelligent thing would be something that anticipates me or my needs in some way. But aside from that I don't really have any profound thoughts on intelligence because it is such a complicated topic and I am really not very well read in that area so I wouldn't want to go out on a limb and say anything more.

LH: *Indeed, it's a really hard question. Every time I ask about intelligence, people have so many different answers. Usually, we see intelligence as something that leads us to build this type of intelligence. In the Western world, we are always confusing intelligence with knowledge.*

DS: Yes. I have another theory that is related to this book I read by David Noble, a Marxist historian. Do you know *A World Without Women?* You might want to read that. It's about the Christian clerical culture of science and how women are deliberately excluded. And what occurs to me is that some of the AI research is actually driven by male fantasies of control and by a sort of desire to replace the woman with some kind of robot. You know, some kind of perfect robot? I have no proof for that but I know that there is something there. This desire to create a super-being is a way of making this more perfect being or companion, right?

LH: *Yes, I've heard about that. A cyberfeminist saying that it is weirdly driven by males because they can't produce life by themselves.*

DS: Oh, yes, that's one of the arguments I've heard. It's a way for men to "give birth."

LH: *Exactly, and also taking control of it. Even if you can give birth you cannot control what's happening, you are not pushing a button saying, "okay make bones now" you know.*

DS: It's much better than a baby, they can just walk away and forget about it for a day! So there is something, I don't know if you'd call it sinister, but there is something uncomfortable about it somewhere.

LH: *I was also wondering, with Watson and Deep Blue, we were almost waiting for the machine to beat us, as if we were demanding this effect. So the question was, do you think that we are wishing to be beaten or wishing to be taken over by something stronger than us?*

DS: I think it's a cultural thing, some societies might be more docile and willing … but clearly AI would fit in very well with any kind of totalitarian society. It would be an awesome instrument of control. As far as Americans go, I can't see Americans putting up with robotic masters. They just feel more independent than that. I mean they hate the government already, so they're more like, "Just leave me alone and I don't want no stinkin' robot telling me what to do," (laughs). On the other hand, the majority of Americans seem to have accepted the institutionalized spying their own government is routinely conducting on their personal communications.

LH: *According to you what's the main cliché, the biggest misbeliefs around AI these days?*

DS: Recently there has been some fear expressed by very prominent public figures, notably Bill Gates, Stephen Hawking, and Elon Musk about the coming age of AI, like large-scale AI, highly advanced. I'm not sure I agree. Something in me always has faith that humans will push back. But I could be completely deluded. We are becoming a surveillance society whether we want to or not. It's coming from the bottom up with cell phone cameras, which most people did not see coming ten years ago.

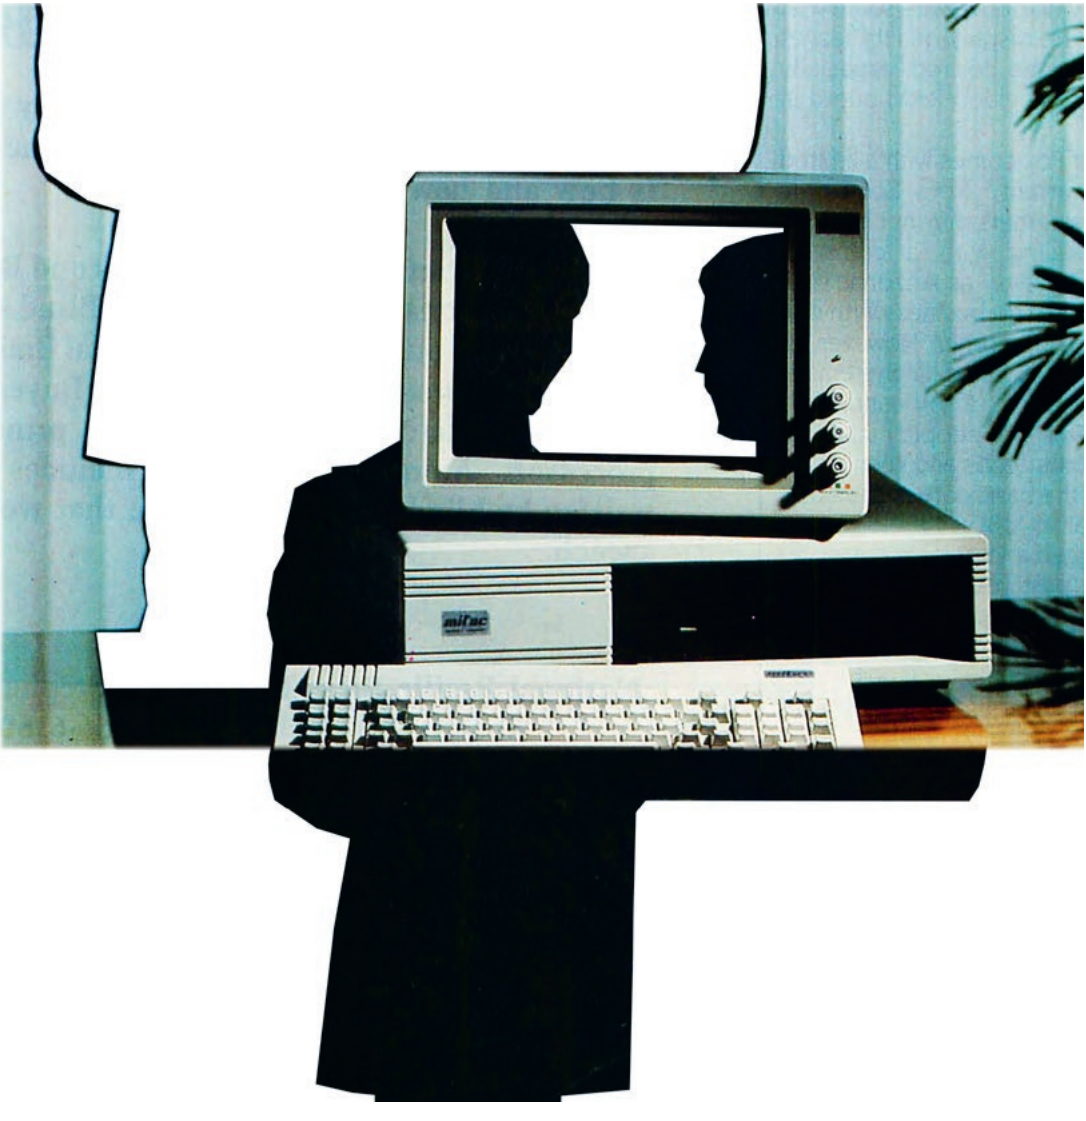

The conversation

A few years ago, we had an exhibit on the history of computer chess and had part of IBM's Deep Blue computer on display. We had a cartoon by Gary Larson beside it that showed a power switch with the caption "Man defeats Deep Blue in one move," and the power switch is set to "off" (laughs). And one of the designers of Watson actually told me, "Hey, you know people think that Watson is intelligent but if the server room that it's in bursts into fire, it would just sit there and burn." So he said, "that's not intelligent."

LH: *What do you think the future will be like technology-wise? Do you have specific visions or even funny stuff?*

DS: I'm an historian so I live in the past! I don't think about the future. It's pretty hard. But you can take a guess and look at Moore's law and make guesses about how many transistors will be on a chip ten years from now, that's pretty easy to do. But what products these magical chips are going to be in is unknown. It's pretty obvious mobile is just going to get bigger and bigger and more important. You still need a good keyboard though so I think we are okay with laptops/desktops in many business environments.

The only thing—if I put on my speculative hat—might be an alternative form of transistor, either biological or nanotubes or something like that. Moore's law just keeps getting refined, more and more, people don't give up on it. We are now down to seven nanometers in the laboratory, at IBM and Hewlett Packard. That's a hundred times smaller than a virus, incomprehensibly small. They can only make that gap so much smaller. Already they are having problems with leakage, which means when you make a little channel for the electrons to go down, when the channel gets too narrow, electrons just start falling off. It might stop here at seven nanometers. If they do stop there, then people

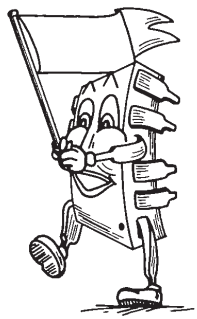

will start thinking about alternatives like using bacteria or nanotubes or something different. Because IC manufacturers peaked in terms of clock speed for microprocessors about fifteen years ago, what they did was to put more cores in them. Unfortunately that created a crisis in the software world where nobody knew how to write code for these new "multicor" processors. They still don't. The single biggest challenge facing computer science today is how to write software for highly parallel or multicore microprocessors because basically the manufacturers, like at Intel and AMD, said, "We can't make these any faster, you software people will have to carry the burden going forward." They hit a wall for speed because the faster it goes the more heat it generates. And in fact, it's a squared relationship, so if you double the speed you quadruple the power that is consumed. One of the Intel microprocessors consumed 150 watts. That's the same as a high-intensity light bulb and it's inside a metal case with a huge heat sink on it and a fan. It's obviously useless for mobile devices. The game now is to make devices super low power and networked. Those are the two themes in electrical engineering today—super low power and mobile in some way. That is what everyone is optimizing for.

LH: *Last question, just out of curiosity … are you a science fiction fan?*

DS: I wouldn't say I am a passionate science fiction fan but when I was younger I was really into it, especially Isaac Asimov and Harlan Ellison, those kinds of writers. Now I don't have a lot of time for reading sci-fi and also many of the things that were science fiction when I was reading or even earlier in the forties, fifties, and sixties is no longer fiction, it's reality. The Dick Tracy watch for example—having a little wrist watch that tells you the time and traffic information, and your heart rate and all this kind of stuff. Cell phones are magical. They are like the communicators

in *Star Trek* in a way, and even Bluetooth earpieces look like a *Star Trek* communicator. Everyone has computers in their pockets now, so all of these things were pure science fiction fifty years ago. They were inconceivable. No one, not even the best futurist could have predicted what we have today. Everyone missed the Internet for example. Microsoft almost missed it completely and Bill Gates wrote a famous memo late in the game, in 1995, saying, "You know this Internet thing, I think it's going to be successful. We better start building products to deal with the Internet." Things move so quickly in our technical field that I don't need science fiction, because I am living it every day!

Mixed feelings V

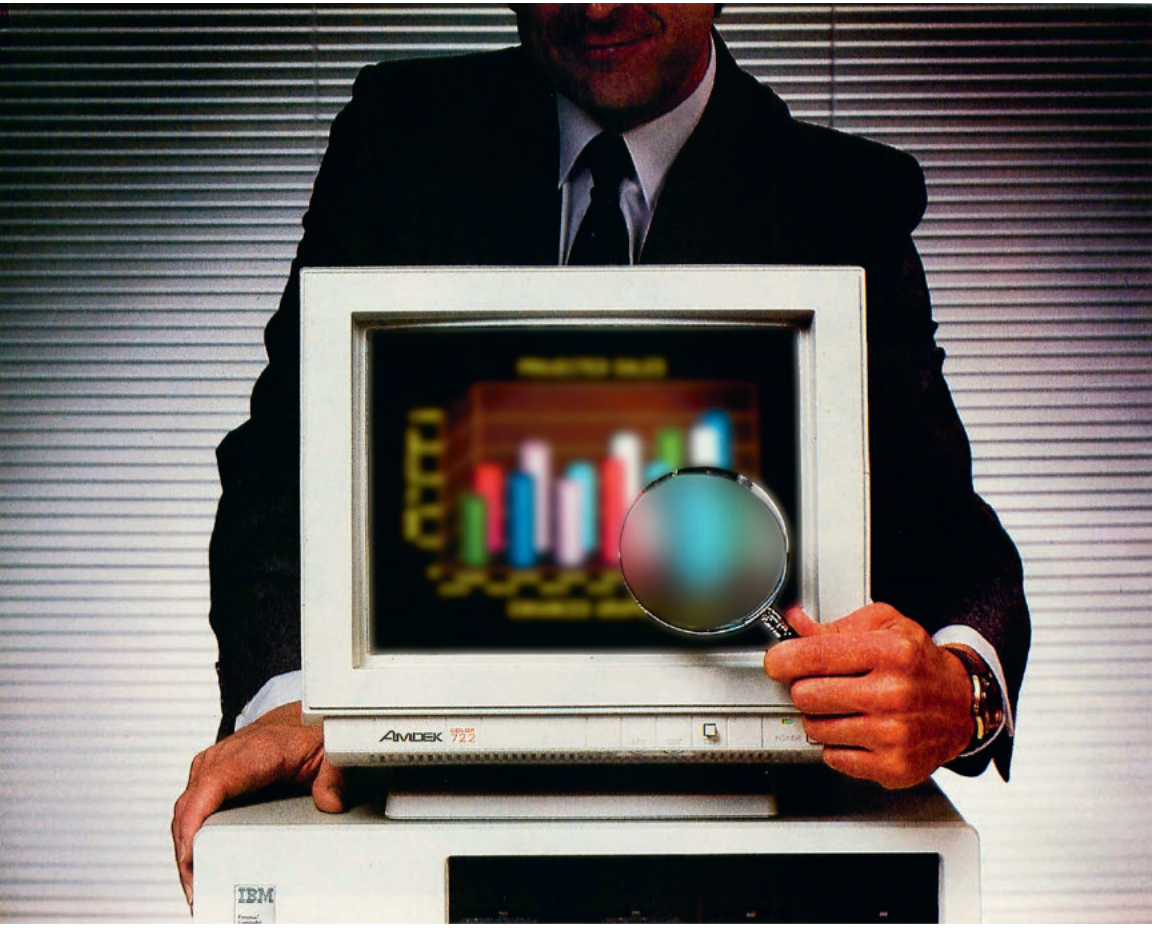

Clear information

DAG SPICER

Dag Spicer is senior curator at the Computer History Museum and leads the Museum's strategy relating to its collection of computer artifacts, films, documents, software and ephemera—the largest collection of computers and related materials in the world.

He also undertakes research for legal and commercial intellectual property specialists and writes on computer history for various media and scholarly organizations. Since he began in 1996, he has given hundreds of interviews with major news organizations and is recognized internationally as a subject matter expert in computing history.

Spicer is on the Editorial Board of the IEEE Annals for the History of Computing and is a member of the American Historical Association (AHA), the Society for the History of Technology (SHOT), and the American Association for the History of Medicine.

INTERVIEW
ERIK DAVIS
WE NEVER MEET THE REAL MACHINE, WE ALWAYS MEET AN INTERFACE

Lauren Huret: *In your seminal book* TechGnosis *you wrote about the intersection of technology and the religious imagination. How did your interest in a philosophical viewpoint on technology and computers in particular start?*

Erik Davis: What would become *TechGnosis* had already been started at university. I was already writing about technology. I studied English at Yale in the mid-1980s, which was the epicenter for post-structuralism and deconstruction in the United States. I was very interested in Jean Baudrillard. I remember very well, I was at the art school at a party and I was talking to this one guy, with a leather jacket, and he goes, "Do you know this book?" And he pulled out *Simulations* by Baudrillard like it was some kind of strange drug or a secret pamphlet or something. So I read that and found it very inspirational. At the same time I was reading early gnostic literature, which I discovered partly through school and just fell in love with. I also started to read the science fiction writer Philip K. Dick, and wrote on him for my senior project. He brought all these things together, a post-structuralist critique of media and the problem of the simulacra with a gnostic notion of liberation from some kind of illusionary world. Very Californian as well. Later when I became a journalist, I continued looking at things in this vein, including the emerging virtual reality scene and what would become Internet culture in the 1990s. I started paying attention to it in 1990–1991. *Mondo 2000* magazine and the cross between the counterculture, the Internet, mysticism, and technology was all part of that mix.

LH: *In* TechGnosis *you wrote, "We associate intelligence with what reads and writes, and nowadays everything electronic reads and writes." I am wondering what you think about intelligence now, how your thoughts on it might have changed and how you would define it.*

ED: I am increasingly persuaded by an argument against strong AI that is similar to Searle's "Chinese Room argument,"

where a non-Chinese speaker writes Chinese using a manual written in his own language and the people who get the results think he knows Chinese. We can look at a machine that is doing certain processes and a great deal of valuable material is coming out for us, so we recognize it as being intelligent. Yet it seems like there is a gap between the ability to program an operation with symbolic units and to have the understanding level that would make that actually a kind of intelligence. I am skeptical about claims for artificial intelligence and even what it means for a machine to take in data and produce what seems to be intelligent responses to it. At the same time, it seems that there are going to be new others in our midst.

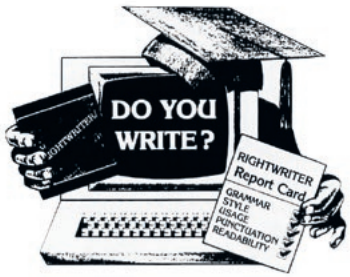

Part of the work in *TechGnosis* is looking at the history of technologies through the lens of fantasy. There is the great line by a political leader in the nineteenth century who is shown a telephone for the first time. He picks it up and he goes, "It talks!" Now we don't think that way anymore, we are just talking on the phone. But in a way, that perception, which you could think of as a very archaic perception, is always part of the experience as well. It's rooted very deeply almost in instinct, in just the very basic responses we have—when something reads and writes it's intelligent. When something talks, it is talking. So there is an element of the uncanny that is now distributed everywhere and that itself necessarily becomes part of the conversation. That is the part I think people miss. And that's *before* we are able to really wrestle with the question of whether or not a given machine is sentient.

Let's say that is our question: Does it pass the Turing test? Does it not? What does that actually mean? How do we deal with the Chinese Room problem? But before we get to answering those questions, there are going to be agents, bodies, corporations, scientists, companies that want to

Trash hand

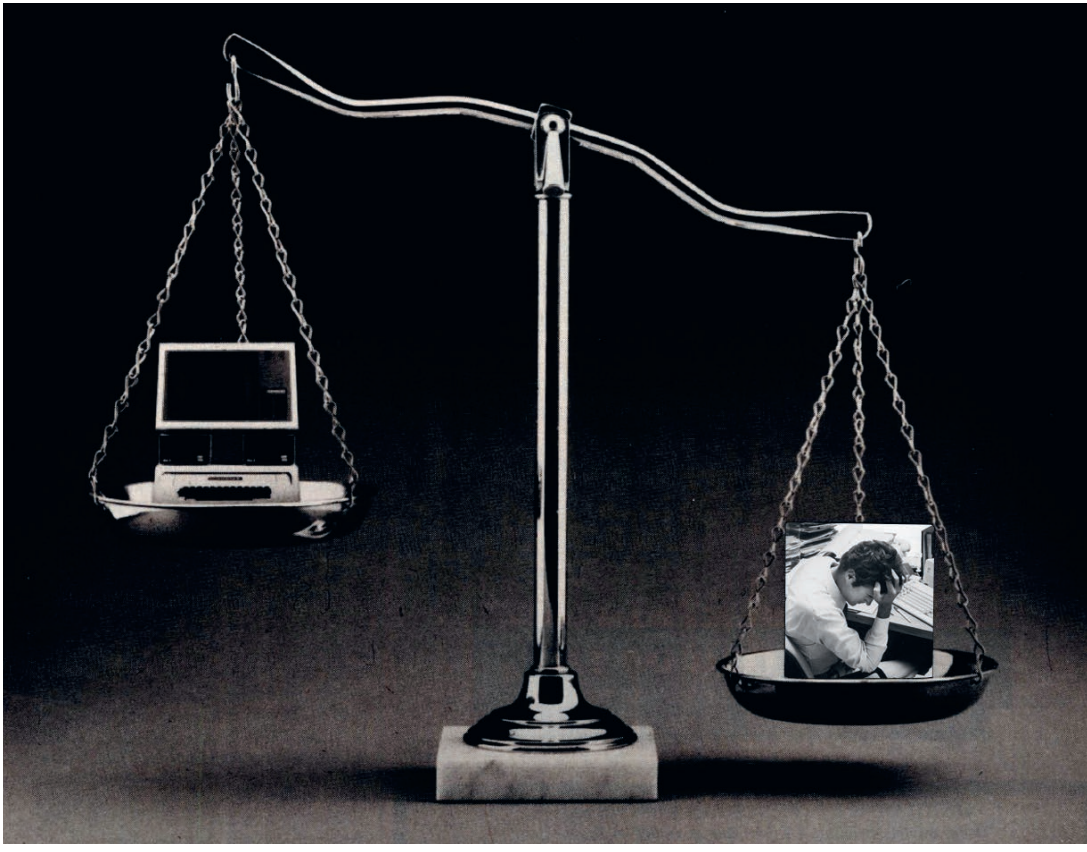

The balance of nature

sell us the idea that we already have intelligent devices, maybe like Watson. And when they do that, they are going to exploit these non-rational parts and stories and myths that we already have, that are widely distributed, that are in our imaginations and our dreams even as skeptical/rational people. We are already in this world of questioning what is real and what is not real. Is it a trickster? Is it a god? All those kinds of seemingly adolescent questions are totally part of the mix because those manipulating populations—corporations, marketers, entertainment companies, and technocrats—are going to use that material. They are already encouraging it and saying at the same time, "Well no it's not really there. We are the ones who actually understand whether or not it is actually intelligent." But in the meantime, we are going to play in this world of animated, uncanny objects and robots that may or may not be intelligent.

Then you have the problem of a drone or a robotic interface that most of the time runs on algorithms and then occasionally gets taken over by a distant human operator and we, encountering those objects, will not know in which of those states it is. So even if we don't believe that machine intelligence is truly sentient, even if we resist that perception, we are still going to be facing this uncanny, half-aliveness in our total environment. It is this weird process where, on the one hand, we are plunging so far into seeing how much we can automate human cognition and produce things autonomously out of that automation—that have effects, that make decisions of some manner—and at the same time, there is this strange, archaic, uncanny revival that is both going to be used and simultaneously attempted to be controlled.

LH: *Do you have an example in mind?*

ED: There is this advertisement for the Charles Schwab brokerage and banking company. It is like a blue machine, it's blinky and futuristic, a kind of shiny monolith, and it says, "I am your investment advisor." It is selling the intelligence of the programmed investment advisor over the human touch, as a plus. After all, humans make mistakes with investments all the time. That isn't going away and we never meet the real machine, we always meet an interface. Of course there are always layers of interface—the machine language, the coding languages that deal with the machine language, the surface screen with the graphic user-interface, the modules, and voices. Then there is the social story that we are told about that machine: "Think different. It's an i-device. It's going to be more creative." In all of these layers, there are elements of spirituality, imagination, mythology, religion, that are going to be part of the story whether you want them there or not.

LH: *There seems to be a real surge of artificial intelligence start-ups in the last few years, in all kind of different fields. According to you, what is driving this recent re-interest in AI?*

ED: I am not really sure even what kind of artificial machine learning people are thinking about a lot of the time. Is it just computationally driven? Does it have to do with bottom-up architectures achieving a certain kind of sophistication? Are they still trying to do top-down AI, logical programming? Is it just associations drawn from big data without any understanding of the underlying principles? In popular culture, it's always been there, these AIs. But it's obvious that we are all wrestling with this issue because some aspect of, I won't call it intelligence, some aspect of computation, is increasingly distributed through our environment. We are constantly interacting in ways where it is very obvious that we don't hold the agency and we are reliant on these devices.

Let me give you an example. One of the things I both love and hate about smartphones are the mapping functions. The thing I hate about the mapping function: With a map on paper, I am able to achieve a large-scale bird's-eye-view of a total space within which I am able to drill down and understand the exact road that I need to get from A to B. I have a larger view with all of the cardinal directions, rooted in my physical relationship with the north pole. And the way that my brain works, that works very easily, so much so that I can often look at a map and then not look at it again as I get where I want to go navigating through space because I have a large-scale picture. Shift to the device, you are immersed in the field, you are like a rat in a maze being told to turn right or left or go forward and it is kind of impossible with the size of the screen and the interface to get a scale-view and have any cardinal objective external reference. So you are almost being held, or you are inside a vector that you are not controlling and you are encouraged to accept a certain kind of passivity that doesn't even necessarily teach you about the landscape that you are moving through. For me that is a real profound allegory for what we are doing. And inevitably we are going to wonder and fear and fantasize about these other intelligences. So it makes sense that it is happening in popular culture. I think one of the most interesting movies here is *Her* because the AI doesn't really give a shit about humans. That is what is cool about it. It enjoys, it will play but that moment that he realizes that she is carrying on a gazillion conversations at the same time and then finally at the end when they are like, "We are just going to move on here. We're done." It's not like, "We're gonna get you! We're gonna get the humans!" They don't care anymore. So for me that shows a maturing in the mythology around AI that was different than Skynet in *Terminator.*

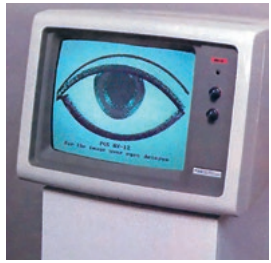

LH: *Do you think the paranoia that machines are taking over humans is changing?*

ED: I completely understand and appreciate paranoid visions around these things. I am immensely distrustful of the powers that be, how technology concentrates power, the new forms of coercion, and all of the political and economical questions that surround those. Whether or not these popular conspiratorial stories are helpful or hurtful or whatever, they are often more than adequate allegories for a more general sense of insecurity about the forms the world is taking.

One of the things that happens too is that the quaint philosophical problems like, "Oh is it actually sentient? Can we really say that it knows? Can we really say that it understands?" will be steamrolled by the sheer power that machine-learning already has. So whether or not we solve that philosophical problem, whether or not it adequately passes the Turing test, in a way it doesn't really matter. And that supports those people who ideologically, or for reasons of power, ascribe to the machine model of humanity—we are just machines, our nervous system is just a neural-net that produces effects. They are happy to go along with that paradigm and just intensify it as much as possible as a way to steamroll over the old humanist analog leftovers. The law of unintended consequences still doesn't seem to have penetrated people's minds. Or maybe it only penetrated the minds of the people who don't have any power. The less power you have, the more sensitive you are to these issues. The more power you have, the more easy it is to believe that we are going to solve these second-order problems.

LH: *What or who can you identify as some of the deep-rooted impetuses and motivations for research and development of AI and Machine Learning? Be it money, control, transcendence, subconscious, governments or other, and do you think motivations have changed recently compared to say fifty years ago?*

ED: There is a shift away from bedrock notions of business, that what you want to do is provide a kind of basis for something, a kind of foundation. Instead now you recognize that once you disrupt an existing part of the economy—like the way people consume media, or concepts about personhood, or about how to get from A to B, all these things—then once these things are disrupted, then there is a space for novel reprogramming. So it is less about efficiently refining the system than it is about building ways of taking advantage of the turbulence that in some sense the technology itself already produces. So there is more room for creativity but the creativity is much more intimately wound up with the trauma of capital and technological crisis.

Meanwhile, there has been this profound shift in what personal creativity means. Creativity has been moved from a zone of aesthetics to industry and it is now thoroughly interfused with capital and technological development—often in ways that, to my eyes, make it much less creative! Such that, a lot of young people, who are also technically competent and very intelligent, experience and desire creativity not to supplement their work but to be their work. It is almost like the capitalist super ego now is: “Create!” no longer: “Obey!” It feels like the big games are on now. And the big games are: installing novel platforms in the spaces that have been opened up through disruption and the kinds of chaos that are going to come. The contest is to see who is going to be able to install those platforms and then adequately capture enough population and capital flow to succeed. I know people who are working on ecological solutions. Their question is, “How can we disrupt the energy market so that we can install this platform that we think is actually

healthy?" They have reasons, maybe they're right, maybe they're not, but they are still thinking in that way. Then there are much more nihilistic forms of the same operation. So I think for young people, they don't get the structural historical turbulent dimension of it. They see the opportunity for personal creativity as well as making a mark and of course making a lot of money. I think that is what has shifted to some degree.

LH: *What happens to counterculture and its relations to technology in this shift?*

ED: The counterculture has a long and complicated influence on the Bay Area technology culture and I think we have reached a point of maximum saturation and minimal conflict or critique. There was, of course, critique in *counter*-culture. Now you have creativity instead of critique. You could look at Burning Man as an example. It was very countercultural, there were guns and terrible people and punks and it was ugly and then, over the years, it becomes this sort of festival of excess capital and libertarian hedonism. Another example is drugs. In the 1960s and 1970s psychedelics had a countercultural character. They were seen as inevitably and necessarily de-connecting you from your conventional personality and social roles. Even if they led to much silliness, they had revolutionary potential. Even if it didn't turn into the revolution, there was a sense of a critical space. How is reality constructed? Who is dominating it? Now it is much more about, again, creating some possible world or playing with the dials on human consciousness to create more gain and bring up the saturation and change the temporality so you get these effects. There are counter examples where people still take them in a countercultural way, in an environmentalist way that is critical of capital or whatever, but on a technical level, drugs have also been incorporated into the creativity industry. Just look at the rhetoric surrounding VR.

Ghost computing

LH: *Can you elaborate a little bit on the relationship between magic and the creation of machines that can learn?*

ED: There is a way to establish an analogy between magic and, let's call it code, because in both situations you are essentially creating a system of artificial or invented tokens that in their interrelationships make certain things happen, at least within the domain established by that code or artificial language. What is the reality of a virtual reality? What is the ontology of some object that is coded that appears only in coded space? I don't know. And I have similar questions about what happens inside of magical environments, where there are certain things that happen that seem to point to a certain kind of consistency. But whether or not you buy that, we are talking about worlds where language and gesture make things happen directly. So in that sense, there is a fundamental resonance that other authors have picked up on, like the Singularity science fiction guy Vernor Vinge in his prophetic novella *True Names.* It's very early, published in 1981, even before William Gibson's cyberpunk novel *Neuromancer* (1984). It is a vision of a collective networked world, a magical environment where hackers who are resisting the government go. It's all built up on a networked code but it takes the form of a visual environment where there are spells, castles, and wizards and stuff like that.

The whole thing about analogy in computers is really interesting because the code is already an analogy of what is actually happening at the level of machine language. When you get down to the machine level, there are a bunch of gates of ones and zeros, of yeses and noes. Actually that is already an abstraction of what is happening, which is just a series of analog voltage spectra that get organized in different ways to register as "yes" and "no" or "one" and "zero". There are ways in which, just from the very basic level, we are already abstracting and telling a story about what is happening on the level of electrons in order to create this whole weird world of computers.

Extra dimension

Enlightenment of darkness

Beyond that, if you look historically and you look at high magic, learned magic in the Middle Ages and the Renaissance, it's almost like reading cyberpunk literature. These guys are imagining the future of information machines. They have this sense that by getting down to the root code, you can code the universe with a set of manipulable symbols. Or you can know everything there is to know, or you can send a message from one place to somewhere else instantaneously. All of these dreams of knowledge, these phantasmic desires of the mind, still drive and reappear and play out in the coding environments that we have now. So there is a historical resonance there, including—to bring it back around to AI—the whole question of other intelligences. So what John Dee was doing in sixteenth-century England, was using an invented language or a channeled language in order to communicate with incorporeal intelligences whose origins and motivations he doesn't really understand. They say things he doesn't know whether to trust. Is it a demon? Is it an angel? It says it is an angel, what do I know? Today, when some agent pops up on a screen and says, "Oh yes, I'm afraid I have to tell you that your flight has been canceled. Would you like to rebook with this other deal that I offer?" I'm like, "What? Who are you? Where did you come from?" We are already in the zone of testing, of not knowing, of secrecy. That is a big part of this esoteric resonance, secrecy. Magic, especially learned magic is very much about secrets. Who has the codes? Who has the information? How is it encrypted? Even if the content of the secrets are very different, the political economy of the secrets are totally different, there is something about secrecy, as such, that shapes reality. And now we are in a creepy realm of profound secrets, the masks of spin, and surveillance. So there is a kind of occult or esoteric quality even without all of the computer game magic, or cheesy surface of pop occulture. Our world is occulted, knowledge is occulted, even as it seemingly becomes so available online.

The real world that we have to make decisions in has all of these hidden forces in it and so the analogies are very strong. From a rational point of view you can say, "Those are just analogies. They don't really mean anything. That is an incoherent way of thinking. We need to think rationally." But I don't think that view is the dominant one. I think the resonances and the stories feed back and come on more and more strong.

LH: *The critical French computer scientist Gérard Berry says that "computers are the most stupid machines ever, very far away from something we can call intelligent." It seems most people consider a computer very impressive, quasi-magical, because it can do a lot of calculation in very short time, something we cannot do. So my question is, if we are fascinated when computers are beating us—Watson is winning Jeopardy, Deep Blue is beating Gary Kasparov—are we actually demanding this effect when computers are surpassing our cognitive abilities? As if to prove or disprove the obsession that humans have with their own power?*

ED: You could have a fun psychoanalytic conversation about that. That secret desire would imply that we are interested in undermining human agency or in checking human agency. There is some part of us that perversely enjoys that. That seems to be true because the dominant story is still one of human agency. It is still human direction. We are the ones programming these things and we are making decisions. We all realize that that is not entirely true and yet, in a way, our culture doesn't really have a good way of acknowledging the ways in which we are not in control, that we are not really agents of our own action. Then, with these computers winning games, where it is really clear we have lost the line, we're beaten. We want Kasparov to win and yet there is something exciting about that loss of control because on some level we know that we've already lost it.

At the same time, I could also say that who we are is not simply cognition or computation. One of the biggest political fights now in this domain is that, through cognitive science, through certain kinds of reductionist neuroscience, through certain computer metaphors taken too far, the computationalist paradigm is quite dominant. The idea is that we are just machines, we process representations, computers process representations, that this is all that we are. Maybe that is true, maybe it's not, and I certainly think it's not true. But one thing that happens is that as we offload cognitive functions into machines, memory, computation, etc., whatever is left becomes more visible. I think there will always be this reminder, though it may grow more surreal and uncanny. There is a huge amount of cognition that we can offload but what remains is not clear. That is also part of the drive you speak of: we want to get the obvious stuff out of the way to get down to the nub of it. Being human is not just about playing chess.

LH: *A computer scientist told me that the more we are advancing in building computers, the less we know about the old codes. He called the codes that we don't change anymore the DNA of the computer and the software. The oldest layers of code are something that remains and something that we cannot take out.*

ED: Right. We are running on all these legacy systems where we are just building layers upon layers. You might put an object-oriented programming language on top of these older kinds of layers and then if you have to go back and fix a problem, you need to get all these old guys and half of them are dead, some can't remember how to do it anymore. So we are spilling forward in this inertial way. To remain conscious of the historical process of construction that goes

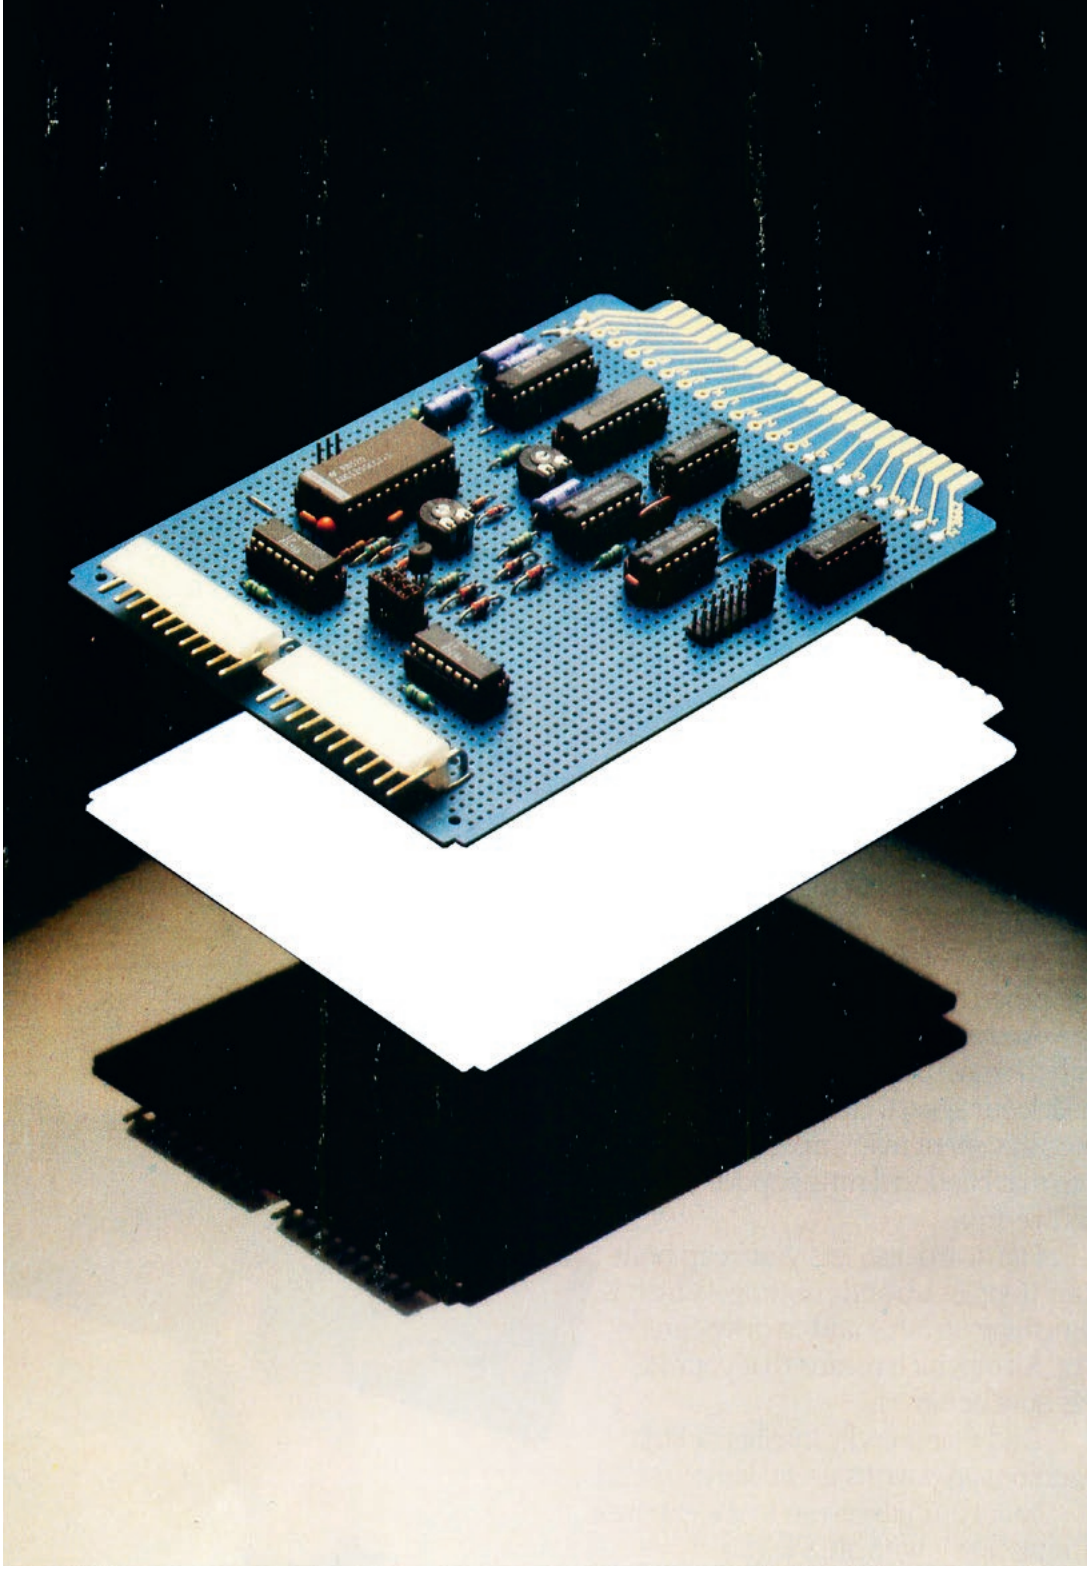

Shadow, void, ideas

into producing these layers is obviously key. We always have to re-embed these machines in these kinds of historical contexts because that allows us to critique both the idealistic fantasies that people have about them and also the strong cognitive computational claims that are made about how we can do this and that. It doesn't acknowledge how much error, ignorance, fuzzy thinking, inertial drag, kludges, and delicate operations are going on at the same time.

LH: *I want to hear about your predictions for the future, considering what we just talked about—this idea of implemented magic in objects and the future of AI.*

ED: There are so many dimensions of the problem and the situation. We could talk all day about technology and automation etc., but at the same time, the real story is global warming and the intense problems that that is going to unleash. I usually quaver at the straight out predictions. It is going to be really confusing! It is going to be really complicated! It is going to be really shitty for a lot of people. But in terms of what we are going to see with AI, I don't know. Are we going to see little robots everywhere? Are we going to have avatars of ourselves? There are so many things that could happen.

I am terrified and sometimes pretty paranoid about some of the drivers of technological ideology. I was watching a TED Talk about autism from a geneticist, talking about the genetic roots of autism. At some point she gave her riff about how the autistic brain works, which is also the way our brains work and it was a completely bottom-up, neural circuit, no higher integration, just a full machine model. She was talking about the way that autistic people will be able to be brought into human society more easily through rich computer environments. She talked about a kid who doesn't speak so he communicates with an iPad and a vocal translator. Then she noted how even social cues and ways of learning how to behave and to respond to other

people could be improved by having really robust engagements with computers as a way to model these behaviors. Then she was talking about how we could all benefit from that. So you have a notion of what human norms are, you embed it into a technology and you use that to refine and create more efficient ways of being. The way in which the fine grain of our subjective, embodied, interpersonal lives is being attacked, or enhanced, if you will, by computational models and the weird metaphysics that underlies those models or metaphysics—that we are just numbers or neural patterns—that is deeply unnerving, and philosophically unsupportable.

I think that a lot of intellectuals with my sort of confusions are in a tough spot because we are profoundly uncomfortable with the forms that a lot of post-humanism takes but we are too smart to be generic humanists or too critical to just embrace the humanist models, which are going to keep going anyway because it keeps everybody happy. Hollywood works! They like the stories of the hero and everything even though we know that that is all bullshit. It still works. It keeps the thing running. So even if a society is post-human we still have these humanist programs that we're running on the platform, like little VR worlds. Well that is not sufficient, and won't withstand the raging Outside that is climate change. But then how to find a kind of post-humanism that is not simply capitulating to power in the computationalist model? I certainly think we have not seen the end of neo-Luddism. What forms it takes, how radical it becomes, is not clear to me and I am not really sure whether there will be a robust intellectual component or whether it will be merely reactionary. Forms of resistance will become a very interesting work in progress.

Mixed feelings I

ERIK DAVIS

Erik Davis is born and lives in San Francisco. He studied literature and philosophy at Yale University. He is the author of four books: *Nomad Codes: Adventures in Modern Esoterica* (Yeti, 2010), *The Visionary State: A Journey through California's Spiritual Landscape* (Chronicle, 2006), with photographs by Michael Rauner, and the 33 1/3 volume *Led Zeppelin IV* (Continuum, 2005). His first and best-known book remains *TechGnosis: Myth, Magic, and Mysticism in the Age of Information* (Crown, 1998), a cult classic of visionary media studies that has been translated into five languages and recently republished by North Atlantic Press. He has contributed chapters on art, music, technoculture, and contemporary spirituality to over a dozen books, including *Future Matters: the Persistence of Philip K. Dick* (Palgrave), *Sound Unbound: Writings on Contemporary Multimedia and Music Culture* (MIT, 2008), *AfterBurn: Reflections on Burning Man* (University of New Mexico, 2005), *Rave Ascension* (Routledge, 2003), and *Zig Zag Zen* (Chronicle, 2002).

In addition to his many forewords and introductions, Davis has contributed articles and essays to a variety of periodicals, including *Bookforum, Arthur, Artforum, Slate, Salon, Gnosis, Rolling Stone, the LA Weekly, Spin, Wired* and *the Village Voice.*

A vital speaker, Davis has given talks at universities, media art conferences, and festivals around the world. He has taught seminars at the UC Berkeley, UC Davis, the California Institute of Integral Studies, and Rice University, as well as workshops at the New York Open Center and Esalen. He has been interviewed by CNN, NPR, the *New York Times,* and the BBC, and appeared in numerous documentaries, as well as in Craig Baldwin's underground film *Specters of the Spectrum*. He has hosted the podcast *Expanding Mind* on the Progressive Radio Network since 2010, and he has his PhD in Religious Studies at Rice University.

"And in conclusion, I'll only use my exceptional powers for the good of mankind."

"That's a vow all we Vector 3005s make. And it's not one we make lightly.

"After all, being the only product on the market with a Vector 3 terminal, a 5¼" floppy, and a 5¼" Winchester rigid disk drive that provides 5 megabytes of storage is quite a responsibility. It used to take 20 floppies to give you that kind of capacity.

"Our powers don't stop there, however. Each 3005 also comes with a 32-bit error-correcting code—the first time sophisticated IBM-style technology has been available on a small business system. This lets us detect and correct errors, and almost completely eliminates data loss on disks due to dirt, wear, or damage.

"All this makes us pretty awesome, all right. But there's more. When coupled with Vector's MEMORITE III and EXECUPLAN software packages, we give you a 30,000 word dictionary, the ability to create your own phrase library, a teaching manual right on the screen, pass word security, plus a host of other word processing capabilities as well as financial planning, forecasting and basic accounting.

"And we're reliable. Our powers won't diminish, our abilities won't fade, and dedication to mankind won't weaken.

"For more information and your nearest dealer, call Vector at 800-423-5857. In California, call 800-382-3367. Or write to them at 31364 Via Colinas, Westlake Village, CA 91362.

"Thank you all for coming today. And I hope we'll have the chance to do business together in the future."

VECTOR

VECTOR GRAPHIC INC.

COMPUTERS FOR THE ADVANCEMENT OF SOCIETY.

Circle 384 on inquiry card.


SIDEKICK
NOTES
JANUARY
IBM

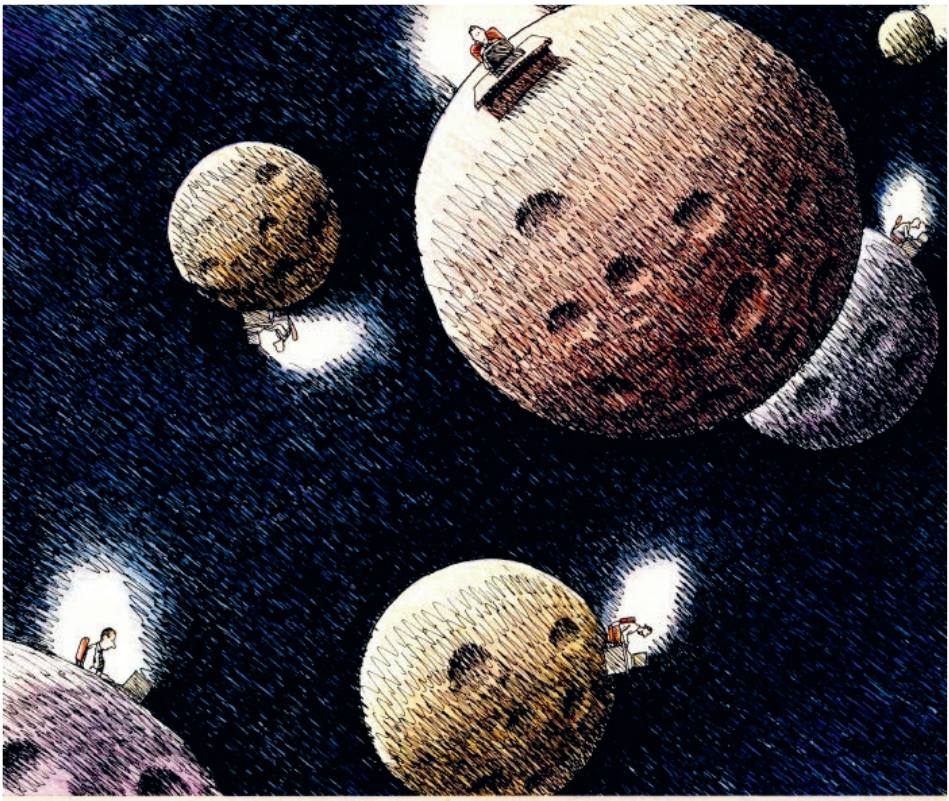

With Network Revelation®, you're not alone.

Span the void that separates you from other PC's. Be as one with a universe of data. Be a true network with Network Revelation.®

Network Revelation is more than a relational database management system. It's a complete applications environment for most microcomputer networks. That's not dreaming about the future. It's low-key raving about a capability of the present.

With Network Revelation, you can send and receive data on local area networks and remote file servers. Rev's data dictionaries let you add or restructure fields at will, saving ages of programming time. And our menu-driven applications generator and procedural language are eons ahead of other databases.

The possibilities are infinite. Distributed processing systems linking worlds. Accounting, inventory and order entry systems connected for instant access to data by a galaxy of users. And your data is secure with complete file or record locking.

REVELATION co-exists with MS-DOS.™ So you can transport Lotus 1-2-3,™ Multiplan® or other data from PC to PC—using popular routes like Novell, Ethernet or 3 Com.™ You can even evolve files from primitive, single-user databases into a powerful Rev application.

Encounter the future of distributed data processing today. Revelation costs just $950.00.* Network Revelation starts at only $1495.00* for a complete four-user system. So call us and we'll arrange for an unforgettable demonstration with a Cosmos rep in your area.

MS™ and Multiplan® of Microsoft Corporation. 1-2-3™ of Lotus Development Corporation. 3 Com™ of 3 Com Corporation.
**Suggested U.S. list price.*

COSMOS™

Cosmos, Inc., 19530 Pacific Highway S. Seattle, WA 98188, 206-824-9942

Inquiry 110

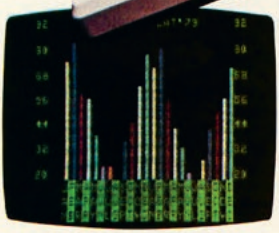

The only computer with color, sound, user programmability and expandability at $599.

The Imagination Machine offers more at its price than any other personal computer on the market today.

Consider these features: 9K RAM, with 14K BASIC in ROM,53-key typewriter keyboard. A fine resolution picture, generated on your television set or monitor in 8 colors!

A built-in, dual-track cassette tape deck with 1500 baud rate, for APF's digitally recorded, "saturated," tape programs. A built-in sound synthesizer. And two, built-in, game style controllers, with joysticks and numeric keypads.

When you want to go beyond APF's library of educational, home-and-personal management or entertainment programs... when you want to create your own programs... you can. The Imagination Machine is programmable in BASIC and 6800 machine language. The Imagination Machine is also expandable. Just add our "Building Block", an optional, four-port expansion device, and you can hook up a printer, telephone modem, and additional memory cartridge or mini-floppy disk drive.

For the name of your nearest Imagination Machine dealer call, TOLL FREE: 1-800-223-1264. (New York residents call: (212) 758-7550) or write: APF Electronics, Inc. 444 Madison Avenue, N.Y., N.Y. 10022.

$599. Manufacturer's suggested retail price.

APF electronics inc.

"YOUR LIFE WILL NEVER BE THE SAME."

See us at NCC, Booth 85

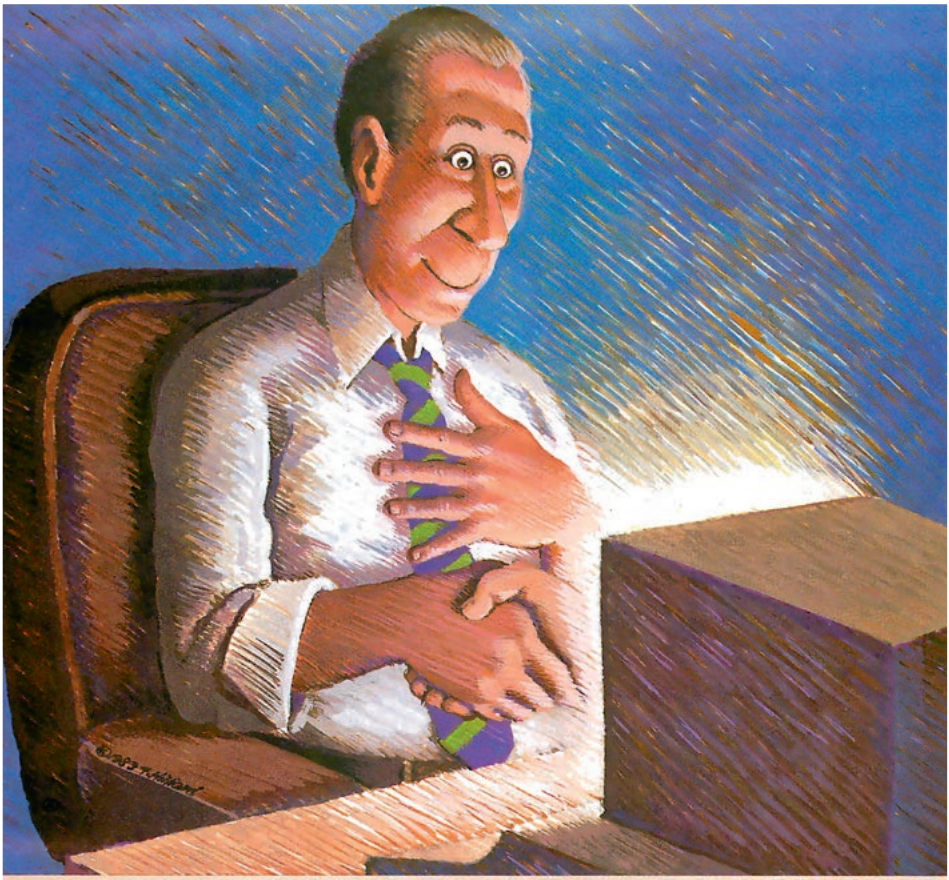

Encounter REVELATION

For $950, Revelation Software by Cosmos will give your PC capabilities that no microcoputer has ever had before. Through a versatile, high performance Relational Data Base Management System developed by Cosmos, Revelation *will give* your microcomputer all the power and sophistication of a minicomputer.

Revelation features a data base with variable-length fields and records, (up to 64K), and unlimited files and accounts. File size is limited only by disk size.

Also featured: A powerful, hybrid programming language.

A fourth generation applications and program generator Our programs write programs!

Communications functions that give your PC the power to utilize applications, programs, and data from thousands of existing mainframe and minicomputers.

8087 math chip support

Revelation co-exists with MS/DOS™ or PC/DOS™, 1.0, 1.1, 2.0 versions, and is compatible with the IBM PC™, XT™, COMPAQ™, Corona™, Columbia Multi-Personal® and the Eagle 1600™. It's also compatible with minicomputers utilizing the PICK Operating System™. Minimum configuration requires 320K memory and 8087 math chip.

MS/DOS™ of Microsoft Corp. IBM PC Registered TM of International Business Machines Corp. COMPAQ™ of COMPAQ Computer Corp. PICK Operating System™ of PICK SYSTEMS. Eagle 1600™ of Eagle Computers. Corona PC™ of Corona Data Systems. Columbia Multi-Personal™ of Columbia Data Products Inc.

COSMOS

P. O. Box AH, Morton, WA 98356

Telephone 1-800-422-2511
Inside Washington 206-496-5974.

Circle 116 on inquiry card.

Someday all terminals will be smart.....

- 128 Functions—software controlled
- 7 x 12 matrix, upper/lower case letters
- 50 to 38,400 baud—selectable
- 82 x 16 or 92 x 22 format—plus graphics
- Printer output port
- "CHERRY" keyboard

CT-82 Intelligent Terminal, assembled and tested. . . . . .$795.00 ppd in Cont. U.S.

SWTPC

SOUTHWEST TECHNICAL PRODUCTS CORPORATION
219 W. RHAPSODY
SAN ANTONIO, TEXAS 78216

Circle 350 on inquiry card.

if He'd used select$_{tm}$ it wouldn't have taken seven days

Learn SELECT* in just 90 minutes. A whole new word processing software concept that kicks the coded key habit and frees you from complicated instruction manuals. *SELECT* is fast. *SELECT* is logical. With single key mnemonics, you'll use dozens of commands that instantly access the rich capabilities of this system. There's nothing like it.

Simply hit "C" and you'll be ready to Create a document.
Key "I" and you'll be in the Insert mode.
Key "M" and Move entire blocks of text . . . and key dozens more.

That's all there is to it. You'll get all that word processing software promises . . . plus a few surprises.

SELECT with SUPERSPELL.* The only microcomputer software with an integrated spelling dictionary. To proof your text all you do, of course, is to key "S". *SUPERSPELL* with its 10,000 word dictionary scans your text at computer speed then displays and corrects all your typing errors. You can increase *SUPERSPELL*'s word power and customize the dictionary by adding new words, one at a time. Ask to see it today at your local dealer.

SELECT with SUPERSPELL . . . just a little byte more.$_{TM}$

SELECT will run on any machine that uses CP/M or MP/M** or its derivatives. It needs 40K of RAM and two disk drives. Special version now available for Radio Shack Mod II*** and Apple II****

* SELECT and SUPERSPELL are trademarks of Select Information Systems Inc.
** CP/M and MP/M are trademarks of Digital Research
*** A trademark of Tandy Corp.
**** A trademark of Apple Computer Inc.

select INFORMATION SYSTEMS 919 Sir Francis Drake Boulevard • Kentfield, California 94904 • (415) 459-4003

Circle 332 on inquiry card.


THE
GREAT ESCAPE!
FROM IVORY TOWERS AND COLD COMPUTER ROOMS TO THE WARMTH OF YOUR OFFICE.
Artificial intelligence for business has arrived in a revolutionary new product . . . Guru.
At last, artificial intelligence designed especially for business! Guru brings together expert system capabilities of artificial intelligence, the productivity of familiar business computing tools and the ease of communicating with your computer using menus, commands or plain English. All available in a single, integrated program.
Guru works like human experts, considering uncertainties, reasoning through forward and backward chaining, asking for more information when needed, and explaining its recommendations.
Guru's expert system works hand-in-hand with all the familiar business computing tools like spreadsheets, statistical analysis, business graphics and a programming language, always available for both expert consultation and your everyday business computing needs.
Best of all, you won't need to learn LISP or PROLOG or buy fancy computers . . . Guru runs on your PC and communicates in plain English! Guru is artificial intelligence that means business.
For more information, call or write Micro Data Base Systems, Inc./Marketing & Sales, P.O. Box 248, Lafayette, IN, 47902, 317/463-2581, Telex 209147 ISE UR.
GURU™
ARTIFICIAL INTELLIGENCE THAT MEANS BUSINESS.
Inquiry 415
mdbs is a registered trademark and Guru is a trademark of Micro Data Base Systems, Inc., P.O. Box 248, Lafayette, IN, 47902
mdbs

The Transporter Night Shift

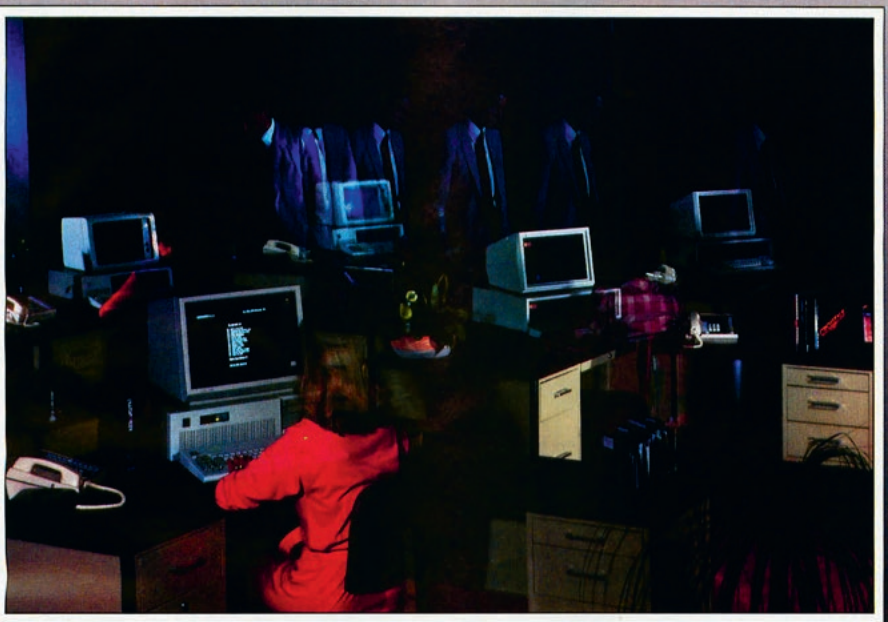

Long After Your Staff Goes Home, We Keep Your PC Making Calls, Sending and Receiving Files, and Keeping a Log

Your microcomputer doesn't go home at five. Why not keep it working?

TRANSPORTER keeps your PC communicating almost as well as your staff could do it. Just load it with a "task list" before you go home. It can automatically call other computers, log in, repeat calls if necessary, send and receive files, and more.

It can do it overnight, when phone rates are lowest.

The next morning, TRANSPORTER gives you a report of what it has accomplished during its night shift.

So you can start the day ahead of the game.

TRANSPORTER comes with our popular CROSSTALK data communications program added on. It communicates with any other system running TRANSPORTER or CROSSTALK. For details, see your retailer or write for a brochure.

MICROSTUF®
1000 Holcomb Woods Parkway
Roswell, Georgia 30076

Inquiry 247

For the IBM Personal Computer. Requires 128K RAM, two disk drives, PC-DOS. Asynchronous Communications interface or equivalent RS232 connection, and a modem compatible with the 'AT' command set.

TRANSPORTER™

interact with a Genius

SUPER XT SYSTEM PLUS®
Ultra-High resolution color card.
Ultra-High (720x400) resolution color monitor.
Gives excellent text resolution.
10x16 character cell.

SUPER XT OFFICE SYSTEM PLUS®
TTL monochrome monitor.
Gives you easy-on-the-eyes viewing.

SUPER XT COLOR SYSTEM PLUS®
High resolution color card and color monitor (640x200).
16 colors.

SUPER XT STARTER SYSTEM PLUS®
Composite monochrome monitor.
Includes color graphics card.
Gives you flexibility to run color and monochrome software.

INTERVIEW
MONICA ANDERSON
THE MIND CAN EASILY BE MODELED IN LESS THAN 10,000 LINES OF CODE

Lauren Huret: *You have been working in the field of artificial intelligence for a long time, as a scientist, programmer and entrepreneur. How did your interest get started?*

Monica Anderson: I've been reading science fiction since my early teens and encountered many stories discussing AI. In college I audited several undergraduate AI classes and was also teaching AI as a teaching assistant. I've been in a dozen start-up companies and most of them had some AI aspect to them. I've mostly been working with human natural languages. I'm a language geek and know about five languages; English is my third. I have also used about thirty programming languages professionally in my career and have created a few small programming languages as needed.

So I've been working with computers analyzing human language using simple techniques like looking for important keywords. It's called "NLP" for Natural Language Processing. I am familiar with all of the traditional ways that computers deal with language and try to extract information from it. Reaching human-level precision with these simple techniques requires quite a lot of effort. After a dissatisfying experience with a commercially sold AI system in 1998, I wanted to leave the field, but at a conference the next year my friend Ted Kaehler gave a presentation that got me interested in trying a different approach. He'd written "Intelligence is survival of the fittest thought" on the whiteboard. That's the principle of Neural Darwinism: In your brain, from millisecond to millisecond, there is a competition between ideas. All ideas are evaluated against your world model to see how well they fit and how good they are. The best ideas get to breed with each other and generate more ideas and some of them are going to be better than their parents. You are starting with random and stupid ideas and they evolve into good ideas after a few runs through this mill. The idea that there is Darwinian competition between ideas in your mind

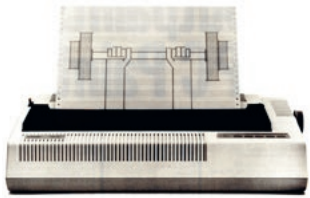

goes back to William James in the 1890s and even to Charles Darwin. I thought, well, I knew about this but never used it—maybe I should actively explore this. I started looking for other kinds of similar AI and after a few months, I had found my way to a more holistic kind of artificial intelligence and I have been researching these ideas since then. It is radically different than the old style, which I call reductionist AI.

I decided I was going to find my own way and think about everything from scratch since so much of the field was "doing everything wrong," so to speak. It turned out that some other people were working on similar ideas. In 1986 a very important book called *Parallel Distributed Processing* was published. Several chapters were written by Geoff Hinton. He published a lot about these kinds of ideas but few people noticed, including me, which meant I missed his breakthrough results published in the mid-2000s.

LH: *This approach to AI led to your own company?*

MA: I worked for Google until 2006 and made about a half a million dollars on stock options. With that money I paid my own salary for the next few years and also employed two other researchers for part of that time. We had a couple of very productive years from 2007 to 2009; we created twenty different experimental systems for this new kind of AI. None of them really clicked but we got partial results from several of them and we figured out what works and what doesn't by experimental programming.

I also worked on a big data project for a company that wanted to create a professional kind of social graph. If two people are on the same board of directors, or if one is the CEO and one is the CTO, they know each other, so by looking at public documents like SEC (US Securities and Exchange Commission) filings and corporate webpages you could make a graph of a lot of people at board level in US companies without having to ask them for the information.

I created a system that could parse information from these various sources and automatically put them in a graph database and connect the dots, so to speak. That was a useful experience and I was wondering if I could create a more generalized and flexible system for others who might have the need to analyze a lot of text. So I started a new company called Sensai—sens.ai—with Michael Gusek and Jonas Lamis in spring of 2013. We set out to provide Watson-level technology at a much lower price point without having to use a lot of big data-style computing power. We did this by providing a subset of the Watson functionality and using some key techniques to lower the computing requirements. We also wanted to make it much easier to use. So I created a new programming language that we call "Content Discovery Language" or CDL, pronounced "cuddle," as the core of a content discovery system. We started out implementing a number of traditional industrial strength primitives and algorithms since you always want to have those as a base. We have now started adding more powerful primitives that use modern machine learning and deep learning methods.

LH: *What is intelligence to you?*

MA: My technical definition is that intelligence is "the ability to perform autonomous reduction." A person looks at the world and uses the process of reduction from this rich world to a simple model for the problem you want to solve. You can reason about models, and you can compute using models, but you cannot compute about the rich world as it is, because it is too complex to enter into the computer. As an example, suppose I wanted a computer to answer a very simple question such as, how many people are in this room? Well, if that's the question, does this pillow matter for the answer? The computer doesn't know. We know. Count the people, we don't care about the pillows or the color of the wall. So the ability to do autonomous reduction is the ability to be able to answer the simple question, what matters?

The leftover (loneliness)

A computer programmed in the traditional way has no clue about what matters. So therefore we have had programmers who know what matters creating models and entering these models into the computer. All programming is like that; a programmer is basically somebody who does reduction all day. They look at the rich world and they make models that they enter into the computer as programs. The programmers are intelligent, but the program is not. And this was true for all old style reductionist AI … If I am ever going to be known for anything, it is Monica's Law, or Anderson's Law: all intelligences are fallible. That is an absolute natural law. There is no such thing as an infallible intelligence ever.

If you want to make an artificial intelligence, the stupid way is to keep doing exactly the same thing. That is a losing proposition for multiple reasons. The most obvious one is that the world is very large, with a lot of things in it, which may matter or not, depending on the situations. Comprehensive models of the world are impossible, even more so if you consider the so-called "frame problem": If you program an AI based on models, the model is obsolete the moment you make it, since the programmer can never keep up with the constant changes of the world evolving. Using such a model to make decisions is inevitably going to output mistakes. The reduction process is basically a scientific approach, building a model and testing it. This is a scientific form of making what some people call intelligence. The problem is not that we are trying to make something scientific, we are trying to make the scientist. We are trying to create a machine that can do the reduction the programmer is doing because nothing else counts as intelligent. But no matter how many programmers you have, if your programmers do the wrong thing, they are not making intelligence, they are just programming. Reductionist AI is just programming. The decisive step is to switch from making models of the world to making models of the mind. That is the key distinction between all the failing AI projects

of the past and the AI of the future. Models of the mind are much smaller. Models of language are enormous and models of the world are gigantic. Nobody has made a model of English, it is too complicated.

LH: *Why "reductionist"?*

MA: I call it reductionist because it is the programmer who does the reduction from rich reality to the simple model that they put in the computer. All of science is like that. Scientists are the ones that create the models and transform, or reduce, them into equations. Reduction is the greatest invention our species has ever made, which works almost everywhere—except in AI. The punch line is, we should be making models of the mind because the mind is smaller than the world.

LH: *Do you have an example in mind?*

MA: If you want to make a self-driving car the standard procedure is to model the car, the road, the pedestrians, the other cars, geography, and then reasoning about these models. The output is data telling you your location and movements—but only within this algorithm space, a logical space. You are not operating in the real world but in the model world, which is much smaller because the color of the sky or the color of the car doesn't matter. So by leaving out irrelevant detail you can reduce the model to a smaller size, permitting you to reason about only that and that is what reasoning is. Reasoning always operates on models. Programmers who create those models are creating the old kind of AI, the reductionist kind of AI.

Hunter Longe: *So what are the other solutions for building an AI?*

MA: First of all, an artificial general intelligence, strong AI, a machine that can perform the same intellectual tasks as a human, has to solve many different problems with the same code. Otherwise it's not a general intelligence. In neural

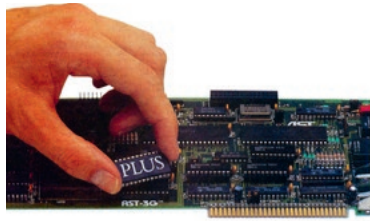

networks, they are what I call model-free systems. Neural networks are what Geoffrey Hinton has been working on. That is the future. That is the technology. Out of hundreds of things that we have tried to make AI work, neural networks are the only one that is actually going to succeed in producing anything interesting. It's not surprising because these networks are a little bit more like the brain. We are not necessarily modeling them after the brain but trying to solve similar problems ends up in a similar design. Neural networks were first invented in the late fifties, maybe earlier. They got very popular, the first neural network conference had tens of thousands of people attending. Then Marvin Minsky and Seymour Papert wrote a paper disclaiming that neural networks are ever going to work. Minsky killed off neural networks with one stroke of the pen. Everybody listened to him, so neural networks never went anywhere.

LH: *Tell me more about the comeback of neural networks these days.*

MA: It started around 2010. The first reason is that the computers got thousands of times bigger and faster. The second one is the emergence of the web with videos, pictures, and countless training material for AIs to learn from. These can just be fed in and the AI will learn from it. The third one was Hinton's discovery how to do things ten million times faster. It took him seventeen years and in the meantime computers had gotten a thousand times faster. So nowadays neural networks are ten billion times faster than in 1986. Hinton was hired by Google and they just released four patents with Hinton's name on them, but which Google now owns. The most important one is called Dropout: If a neuron in the neural network has a hundred incoming connections it will listen to a hundred inputs and ignore twenty

of them at random, then it does something, learns from that and ignores twenty others and learns from that again. This results in multiple possible, hypothetical or near-misses of the thing. That way the training material can be used a thousand times more. Conducting a thousand experiments always ignoring a certain amount of input, means that one can still learn more from the same training material. Then came Hinton's "deep learning," based on these tricks and neural networks. Deep learning is now the number one buzzword in AI.

LH: *And deep learning is not reductionist?*

MA: Deep learning can solve multiple problems. The epicenter of deep learning research was at the University of Toronto, Geoff Hinton's lab with Joshua Bengio and Yann LeCun. The fourth big name is Andrew Ng from Stanford University. These Canadians developed image and video analysis better than anyone before. You can show pictures of things to the computer and it will tell you, "Oh, there is a woman in a white tennis dress holding a tennis racket and behind her is a car." The technologies used by Google today are there because they hired Hinton. Recently Yann LeCun, who has been running the AI department at Facebook since 2013, posted an article called "Text Understanding from Scratch": starting from a system that doesn't know any language and feeding in Amazon.com product reviews, you can train the system to recognize positive and negative reviews. Using five-star reviews and low-star reviews and feeding into the neural network whether the review is either positive or negative in order for it to learn. Then you give it new reviews that it hasn't seen before and it sorts them into good and bad very nicely without knowing what the star-rating is. That is a generalized classification problem, specifically it's sentiment analysis, which is a big deal. A lot of AI today, especially in the social media space, is about sentiment analysis: "What are people thinking about my

Water screensaver

new product? Do they like the Apple watch or not?" You want to have thousands of Twitter feeds to figure out whether the Apple watch is cool or not. That is a very big money maker if you can do it right.

LH: *So deep learning AI can do that?*

MA: Deep learning is not strong AI, but even though there are still some problems, it is the only thing we have that is on the path to real AI. Take restaurant reviews for example, there is always something good and something bad in every review and you can't really say which way it goes. The precision with the AI systems up until now has been 65% to 70% right answers out of 100. That is nothing—50% is a coin toss. But the above-mentioned researchers reached over 90%. They completely out-classed every reductionist system with just a system that learned from examples. It is called supervised learning because the examples were labeled good and bad. They beat the state of the art by over 20% and this happened in other domains too such as video and signal processing. This is not science fiction because even in Android phones the voice recognition is much better than the one in the iPhone, because Apple uses the old kind of AI. The Android uses the new kind and it gets it right much more often than the iPhone does, since it has a much larger vocabulary and is much more tolerant to dialects and linguistic variety. So it is already in products. Google made it into a product and there is a rumor that Google internally uses deep learning already in about sixty places and applies it in their product line. So Google is gung-ho for this and other companies are too. Even though it's almost impossible to hire anybody who knows anything about this at the moment, it is pretty obvious to me at least that deep learning is the closest thing we have to AI and it's a big deal. It's a

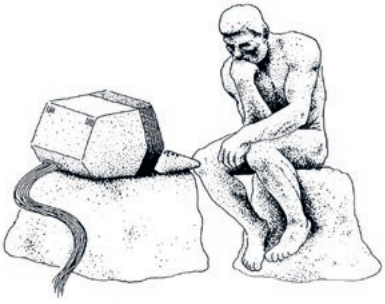

revolution in multiple disciplines, signal processing, video, text, sound, speech, DNA. It is going to be applied in many fields, like robotics, where robots could learn to walk like a child, continuously improving, up to ballet dancing. That is very different from having to program everything. The system just learns from its mistakes.

LH: *Do you think it will be possible to model the mind or even to understand it?*

MA: Everybody thinks that science is the hot thing, but science is a package of methods that we use to stay sane. Humans are basically not scientific. In fact, we are intelligent without being scientific. We have been intelligent for hundreds of thousands of years without having access to science. So science is just a latecomer bunch of methodologies that are used for understanding deeper things and it is mostly for experts. In my talk about Bizarre Systems, I explain that scientific methods fail, even in very obvious situations of everyday life: we find our way across town, buy a newspaper, come home, read and enjoy the newspaper while listening to a symphony on the radio. How can a computer enjoy a symphony? Does it even mean anything? Reading a newspaper is very similar requiring all this recursive structure in language to understand. So a computer or robot or an AI that understands is where we want to go. To avoid the complicated term "machines capable of autonomous reduction" I just say, "understanding machines." You understand the world, rather than reason about it.

LH: *What does it mean to understand?*

MA: There is an extreme distinction between understanding and reasoning that most people don't get. Reasoning is a step-by-step logic that requires perfect data as input, a model you can reason about, which will have an answer as output that can be applied to the world somehow. But understanding is different. Intuition-based understanding means taking

Wow

in the sensory input, knowing what matters, learning what is new and knowing what to do next; understanding the situation, understanding language, understanding what is dangerous in the room. If you meet a saber-toothed tiger, you don't have time to reason about it, you know what to do instinctively. 99.96% of what we do is of that intuition-based understanding, which uses past-life experience and the learning we have done over a lifetime to do what we do on a daily basis. All language use is done without reasoning. You wouldn't have time to reason, "Okay this is plural, is it possessive or not?" You can drive your car without reasoning. Knowing the difference between understanding and reasoning is key to AI. The old kind of AI is all about reasoning but they don't have any understanding underneath it. They need this human capacity, the subconscious understanding. That has to be the base of AI.

LH: *Can you describe an example of the things you are doing?*

MA: I write Java programs. I start the Java program and it starts out empty. Then I feed Jane Austen novels into it and after four pages the program figures out there are words and after a few dozen more pages it recognizes the names of the characters. At the end of the book, I can start asking it comprehension questions. Even though these are very simple questions, one can test if it understands language to some degree and what needs to be adapted and changed to improve it. With holistic systems like this, there is no easy reductionist formula to say what to do next but you need to experiment and look at the output. It is a black box just like any brain is. Learning from mistakes is key. You start out empty, you don't know what matters, but after a while you start noticing what matters because you have something to build on. People think they can teach their kids language, they do, but they do it while speaking to each other. Children learn it from observation mostly, from studying examples and that is probably true for AIs as well.

LH: *It's very complex.*

MA: As I said, there are hundreds of ways to do AI, many of which we tried out. We look at the existing intelligence, human intelligence, or we look at neuroscience to try to figure out how the brain works. But I'm fond of saying, "The story of what the brain needs to do is much shorter than the story of what the brain does". If you want to make an AI, don't look to neuroscience, just model the mind! How do you do that? You can look to epistemology. Epistemology is the discipline that discusses what knowledge is, what learning is, or how we know what is salient. Science doesn't tell you what matters; it has to be answered before you make your equation. Epistemology is at the bottom and science is this thing at the top. Epistemology talks about the limits of science, but there are people who don't believe in the limits of science. They think that science can do everything, solve all of our problems. This is a faith-based exaggerated belief in the power of reductionist science, called scientism. It's a disease. Its about as misleading as religion. The modern superstition is scientism.

LH: *What do you think about popular speculations on AI, especially its dangers?*

MA: Do you mean the idea of an AI take-over, of artificial intelligence getting to be super-intelligent and taking over the world and killing all of the humans? I'm glad you're mentioning that! First of all, the people who think that we are going to have a super-intelligent machine are the ones who believe intelligence is logical and they believe we can have arbitrary, powerful intelligence. But if you consider "the frame problem"—the world changes behind your back etc.—its immediate implication is that all intelligences are fallible since they never have all the input information and are always working with imperfect information. So if we have an uppity AI that wants to take over the world, we just wait until it makes a mistake and then we pull the plug on it.

That is a straight-forward answer and it follows directly from the frame problem theory. There will always be things unknown to you because the world changes. Recursive self-improvement is impossible because AI is so simple and therefore there is no chance that AIs are going to quickly become super-intelligent and take over the world. That was the point I wanted to make. You can talk to Nick Bostrom, Ray Kurzweil, Elon Musk, or Stephen Hawking. They are talking about the dangers of AIs, but they have no clue about these things. They believe in what they read in Bostrom's book and his book is written from the standpoint of someone believing in the old-style reductionist AI who imagined that AIs are going to be these big monstrosities that can self-improve. That is not where we are going.

Deep learning neural networks are a much better alternative and these systems are so small—a few hundred to a few thousand lines of code—that there is almost nothing left to improve! A superhuman, logic-based, infallible, omniscient God-like AI is totally impossible and you can quote me on that.

LH: *What's your vision about the future of AI?*

MA: They are going to be something like an improved personal assistant in a smartphone, like Siri or the equivalents from other vendors. More and more AI programs for various understanding tasks are going to appear. Google will probably have hundreds of deep learning programs and AI in their system and we will never know except our searches are going to be better. We will have generational improvements over time and we will have control over those improvements. If you've seen the movie *Her,* the devices they have there are basically the goal of my AI company Syntience. Starting with language understanding, the aim is to produce something you can carry all day and talk to it and it answers.

These are called Confidantes™. Maybe you have a dozen experts in your Confidante. "I want to talk to John because I want to talk about patent law." Then John your patent expert comes on the phone and talks to you. If you have a question, you ask one of your experts. If they don't know, they will look on the web. If they can't find anything on the web, they will ask other Confidantes on the web. If no other Confidantes know the answer, then they pick some Confidantes to ask their humans. So work in the future is going to be you carrying around your *Her*-like device, your Confidante, all day long. It tells you things that you need to know without you asking for it. How a legal change effects you or a market opportunity or a vacation spot or flight info, or whatever you might be interested in, it knows about it and tells you about it without you going and looking for it. Confidantes will raise your IQ by fifteen points or something like that.

LH: *How do you think all this type of knowledge that we have already will affect our personality and our mind?*

MA: We carry around our entire life in our heads, our education and all the experience we have and we understand how to interact socially because we have practiced it a lot. AI is going to do the same thing and it will be a symbiotic relationship. We are going to have personal AIs that are going to help us and it is going to be a democratic thing. We already benefit from such symbiotic relationships, which are going to be even easier in the future. AIs will have much more knowledge, replacing our own web searches. If we are having a conversation and I use a word you don't understand, then your Confidante will just explain it to you because it knows you don't understand. It will say, "You should buy milk on your way home because you are out and by the way, I did your taxes." That is the outlook, I think. In fifteen years, everybody is going to have their own Confidante in something like a cell phone or implanted.

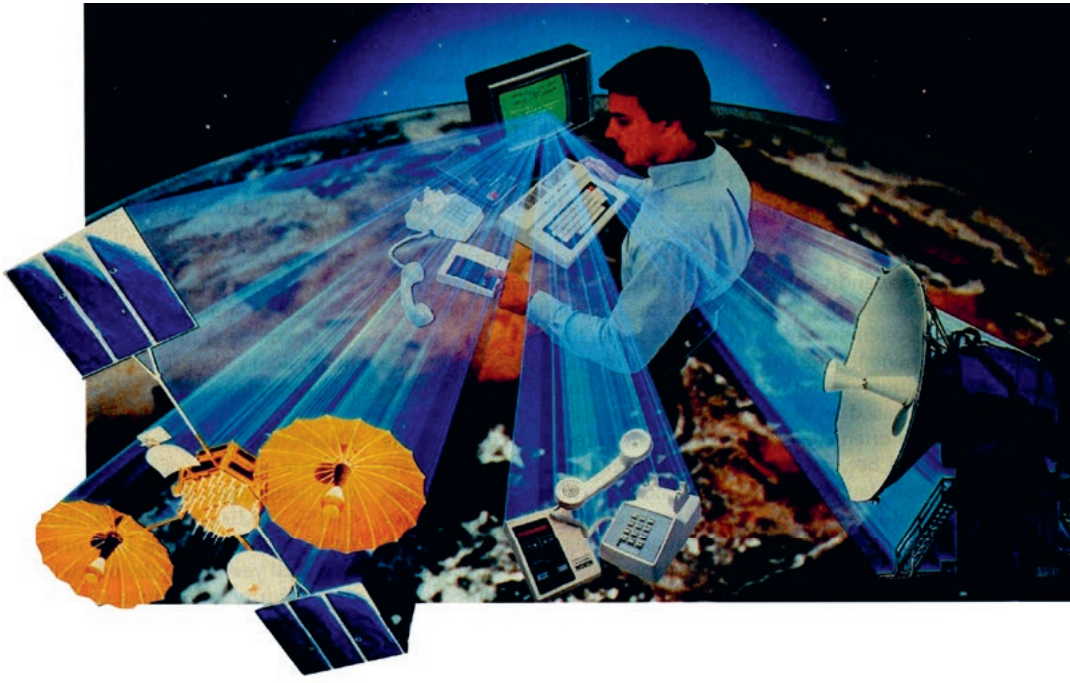

The center of every miracle

I don't really have a problem with that even if I don't believe much in implants. Everybody is going to be smarter. We are going to cooperate more. The Confidantes might be tracking money so we don't have to carry money anymore. Everything is going to be bartered and the Confidantes are going to negotiate the barter transparently to us and lots of overhead falls away. Research goes faster because all the researchers are continuously talking to their Confidantes and they are talking to each other. It is basically a big group of people thinking together mediated by their Confidantes. The world becomes a better place.

LH: *So, having worked in the field for a long time, what can you identify as some of the deep-rooted impetuses and motivations for research and development of AI and machine learning? What are your motivations?*

MA: My deepest motivation is simple. I would like to live for a very long time. I am a cryonicist; I signed up to be frozen after I die. Somebody your age is very likely to live several hundred years minimum, probably thousands of years because we are finding cures for all of the major medical problems we might have. Now, we may be able to push life spans out several decades. If you can extend the life span by ten years, in those ten years we get more medical breakthroughs and so on. I'm at the borderline, I might die, I might not. I think we need AI to solve most of our health problems and most of our political problems. There are people working on the medical aspects, some people work on the political aspects, and if I want to contribute to longevity, I have to work on what I know. Swedes have a saying, "Dig where you stand." I'm digging at AI, hoping to get something that is useful for humanity to solve its big problems and to me longevity is the first one I need solved.

LH: *Can you explain a bit more about cryonics?*

MA: This bracelet I am wearing has instructions for what to do if I should die. There is an organization called Alcor Life Extension Foundation, which is headquartered in Scottsdale, Arizona. They have a few thousand people who signed up to be frozen after they die. I think I was only the 300th person or 299th person to sign up at Alcor, I have a pretty low number. They have two people that were born in the nineteenth century that are frozen and the first person ever frozen, Dr James Bedford, is also at Alcor now. He was frozen by another organization but they went bankrupt. It costs about a thousand dollars a year for the membership and then you have to pay life insurance premiums of a few hundred dollars per year to cover a payout of 140,000 dollars if you want to just freeze your head, because that is all that is important. If you want to freeze your whole body for superstitious reasons, then you will have to pay premiums to cover a payout of 250,000 dollars. If I die—over 90% of all deaths are predictable several days in advance because it is mostly people declining in health—my friends will call this phone number and Alcor will send their specialized medical truck. In the truck they can freeze the person, then drive it to Scottsdale where they have dozens of tanks. Over a hundred people are frozen already.

LH: *How many?*

MA: I think 110 or so, plus a half dozen dogs. If you are signed up to be frozen you can have your pet frozen for a 10,000-dollar pet fee. There are other organizations too, like the American Cryonic Society, which have more subscribers but they are doing a slightly simpler freezing process and they are cheaper for that reason. There are offshoots in England and a few other places. If you are about to die, the best thing you can do is take a plane and go to a Scottsdale hospice specialized in this. You go there, you die, and you get frozen. Their doctors replace your blood with an antifreeze

and then they have a computer-controlled process that cools it down to the temperature of liquid nitrogen. They put a microphone in the skull to listen for cracking, to make sure that the brain isn't cracking into pieces—a major problem that they had to fight. We basically hope that one-day medical technology will progress until we can thaw those patients out and repair whatever is wrong with them, including the death in the first place. A lot of people think cryonics is perfectly reasonable. We know that it might only have a 1 or 2% chance of working, but we know how good the alternatives are! I'd rather pay a few thousand in life insurance premiums total over a lifetime, for a 1% chance of living to be a thousand years old, than have the certainty that I will never come back after I die because I got cremated or buried.

Mixed feelings IV

Focus

MONICA ANDERSON

Monica Anderson is co-founder of sens.ai and Syntience Inc. She has used industrial strength traditional AI methods for natural language processing and expert systems for decades and has worked with AI for Google, Kosmix, Kanisa, Cisco, and Schlumberger. Since 2001 she has researched deep neural networks for understanding human languages and has devised and championed concepts like Model Free Methods, Artificial Intuition, and West Pole (Blank Slate) strategies for Connectome Algorithms and she recommends Epistemology-based approaches to artificial intelligence. Recent advances in AI based on Deep Learning have largely vindicated her theories and AI implementation strategies. Several of her key presentations are available as videos at http://videos.syntience.com.

INTERVIEW R.U. SIRIUS

THERE IS A MECHANICAL ASPECT TO OUR PRESENCE IN THE WORLD

Lauren Huret: *Please tell me what your everyday motivations are to do what you do? What are your strong interests and motivations?*

R.U. Sirius: I go through a dozen different moods on a daily basis ranging from totally not giving a fuck about anything to believing that I could be the world's great mutator of the human condition and all points in between. My motivation for doing what I do is to be a channel, a medium for other peoples' ideas. I like being an interviewer myself. I've done podcasts. I love just tweeting stuff that I find every day when I web surf and sharing it with people. I like bringing some humor to a lot of areas that tend to be very dry and a bit pompous. This makes it easy to be funny because it's always inappropriate.

As a writer and editor of *Mondo 2000* I developed a reputation as being associated with cyberpunk, cyber culture, tech culture, and all that, so I tend to be asked to write or speak about those things. After *Mondo 2000* I have been editor-in-chief of *gettingit.com,* in 1999–2000, *H+* magazine, *Accelerator,* and a few others.

LH: *Tell me what you think about recent breakthroughs in artificial intelligence and how do you think it's implemented in our daily lives, or is what we currently have really AI in your opinion?*

RU: There is this distinction that is made by some AI scientists between artificial intelligence and artificial general intelligence. Artificial general intelligence is the idea of an intelligence that has the flexibility and the ability to respond to diverse inputs, questions, and challenges, perhaps mimicking the human brain or perhaps being something else that is distinct from it but can have some of the similar skills. What we have now is artificial intelligence, which are machines that can handle a couple of things but can't really respond to unexpected questions or stimuli or anything like that. There are developments within commercial AI of

machines that are starting to show more abilities to respond to unexpected stimuli and so forth. I think it is a matter of just more complex algorithms. But it's not really doing anything that mimics the flexibility of a living system, even a cat or a dog, have the invariability to respond intuitively and quickly to stimuli and unexpected actions. There are all these projects to build brains based on the human brain … a whole bunch of them and maybe something will happen with one of those. Of course Ray Kurzweil has become the chief artificial intelligence engineer at Google. So his desire to create a Singularitarian intelligence may help shape the eventual creation of something like that, if it is possible. I'm not enough of a scientist to know whether these things are possible or not. Is it a whole question of computing power becoming able to deal with greater amounts of information than the human brain? The question is whether having information can turn into contextual responses and actions in the actual physical material environment of the world. I think there is a huge gap between this idea that the intelligences of our machines are doubling every year, and actually having something that can respond to the world in its environment. But you never know … it could happen …

LH: *I pulled a little quote from Erik Davis' book* TechGnosis, *that I really like. He wrote, "We associate intelligence with what reads and writes and nowadays everything electronic reads and writes." I am wondering what you think about the way we are seeing intelligence now. Has it changed from the last fifty years regarding the developments of new technologies?*

RU: I think there is a strong tendency within the tech world and within the AI world to equate intelligence with logic—deductive logic—and I think, as Jaron Lanier has been

pointing out, that rather than making artificial intelligences that are as great as human beings, we might be trying to reduce human beings, to cause them to think more like artificial intelligence and engage entirely in cold deductive reasoning. And that can lead to some very weird politics. You have some people just becoming out-right fascist saying, "Well I've done the numbers and the only way humanity will survive is to have a massive authoritarian state. I have looked at the statistics and these people are inferior because of these statistics," and so forth. You begin to see a lot of that now, it's increasing. I think that is, in some ways, human beings wanting to be logic machines. Of course, human intelligence, and even some animal intelligence, lies beneath the surface, though logic is very important. Irrationality in terms of science and politics is a problem but then at the same time there are these other engines that operate within human beings that help us to create culture and help us to create technology. The great geniuses of technology tended to operate in very intuitive ways, like Einstein. And Francis Crick did his part in discovering DNA while using small doses of LSD on almost a daily basis. So a lot of this stuff emerges not by thinking about things, but by having enough hard data and information and then something that's almost mysterious emerges and you have an idea that is supple enough to be true and workable and possibly humane. A lot of what goes on with people is trying to reduce their thinking to the way they imagine artificial intelligence thinking … it tends to create brittle technological systems, making it harder to use. You buy some product and you wonder, "Didn't anyone try to use this like a human being and discover that the most important thing is missing?"

LH: *I found an article where you wrote about the constant data flows in Cronenberg's* Cosmopolis, *especially the quantification of the body.*

RU: 

Quantified life! People running around recording their data all the time so that they can improve themselves. Of course they have the wristwatches, now … they have Fitbit and more complex machines, and that is exactly what was happening in that film. That is another aspect of what I was talking about—human beings as mechanical systems. If it becomes completely intuitive and you don't spend a lot of time thinking about it, then it's great. But if you are living your life to measure everything, it's an accountant's vision of the Singularity.

LH: *Apparently it seems like the tech industry is trying to avoid the effect called the "uncanny valley" at all costs. What do you think about this?*

RU: I think it's an interesting idea. I don't really think that the tech industry is a group of people that all gather together and ask, "What are our goals?" Maybe some of them are trying to avoid that uncanny valley and the possible panic among people and the opposition among people that might come with it. Then there are people out there like Dr. David Hanson, who want to create humanoid intelligence. Have you seen the face that looks a little bit like Aubrey de Grey or Abraham Lincoln, and it will respond to your questions? But if your question is a little tricky, it can't really respond to it. The Japanese particularly like the idea of realistic robots. They want something to keep company with their aging population. When they need treatment, then the person can feel like there is somebody in there. Even here in the West, people like the idea of robotic pets that look and seem like pets.

LH: *When I am feeling the effect of the "uncanny" in front of robots, it is basically reminding me of my own weirdness, how I am in the world moving and talking. It is very awkward to see a machine trying to imitate me.*

In the devil's eyes

RU: There is a mechanical aspect to our presence in the world. I mean the conditions for our existence and biology and so forth are preset. Transhumanists are trying to change that a little bit but they are still dealing with the code that is presented to them and then they're trying to tweak it. So yes, seeing something that is a machine sort of seeming like one of us can be disconcerting because it reminds us that there are aspects of robothood in what we do.

LH: *There seems to be a real surge of AI start ups over the last several years as well as Hollywood films, TV shows, and also an increasing "Internet paranoia." What is driving this recent re-interest in AI?*

RU: People in Hollywood follow and learned their lesson since the early to mid-1990s when they didn't know anything about what was going on in the tech world. The entertainment industry was ignoring it until it totally came up and bit them in the ass. Brett Leonard, who made the film *The Lawnmower Man,* was the guy in Hollywood who was going around and telling people that they were ignoring the next economy, the next culture, the thing that was going on and it took them quite a while to catch on. James Cameron was hip to what was going on. Other than that, it took them quite a while to catch up. And since then I think they follow tech culture and it just makes for great narratives. The fear of the human-like AI's; the hacker as the ambiguous hero or the ambiguous villain. All those things, the idea of somebody bringing down economic systems or corporations by hacking, by being sneaky. These are appealing. I think Bruce Sterling pointed this out, there was always a kind of a Spy vs. Spy thing to the hacker aesthetic and ethic. I don't think things like Chaos Computer Club, WikiLeaks, and Anonymous and all that really could exist without having a relationship with their enemy.

With the absence of the bad guy to rip-off and hack and screw, that sort of thing wouldn't happen. And all that makes great TV and great film because it has that mysterious Spy vs. Spy thing but it also has the potential for really complex, multilevel plots involving data and different ways of thinking. The best thing these days is *Mr. Robot.* I was watching it last night and I thought, well, this is actually better than *Neuromancer!* And it's a TV show and I think people are having a hard time even acknowledging it. That show is almost a gift to Anonymous!

LH: *What are some of the deep roots and links between AI and Transhumanism?*

RU:

In some ways there are two distinct factions or roots in Transhumanism. One is Singularitarianism, which is really focused on artificial intelligence. It's singularly focused on artificial intelligence and the idea of creating machines that are many times smarter and many times more effective than human beings so that it and we can, in turn, resolve all our problems—not die and live in some kind of a utopia. Whether it's a wondrous utopia or an accountant's utopia, I'm not sure. Transhumanism really comes out of the idea of extended lifespan and possible immortality. Transhumanism, certainly when it started, didn't have a strong focus on artificial intelligence. The early days of Transhumanism in the 1970s and '80s, they were mostly interested in genetics and the genome and what can be done to perhaps tweak the genome to allow us to extend our lives, and pharmaceuticals and nutrients and different things that human beings can do to live longer. Then there was also the general technotopian or romantic technological attachment. I guess AI was hot in the 1970s. There was Marvin Minsky with *The Society of Mind.* There were some people who were expecting a

breakthrough right away; that there were going to be really brilliant, not human-like, but really intensely useful AI machines that would fundamentally resolve big human problems like resource scarcity and so forth. And that kind of collapsed. That generation of AI research didn't seem to be going anywhere and people kind of lost interest in it. But I think Transhumanists have a pretty broad spectrum of technologies and ideas that they hope will generate the future they desire—molecular technology—in other words nanotechnology is very important. Genomics is important. I think at this point, most Transhumanists probably would say that they believe we can get a super intelligent AI and that will crack how to most effectively use genomics to extend life, how to most effectively use molecular technology to end resource scarcity and to extend life. At this point AI is probably a pretty big deal in Transhumanism and Singularitarianism both, but there are certainly elements in the Transhumanist world that look less to AI and more toward other types of technology for increasing intelligence—different types of enhancement: cyborgian things that you put in your body or add to your body, anything that can enhance human skills—an ability to jump 100 feet in the air or fly like a bee or whatever, see in the infrared spectrum. Grinders have already come up with these eye drops that let you see in the infrared spectrum at night, they give you night vision. So there is a lot of that stuff in Transhumanism. It is fair to say that Transhumanism has a bigger street-level group of enthusiasts, like the grinders, the sort of cyberpunks, neo-cyberpunks … a lot of people who don't have a lot of money and like to mess around with biotech in their garages and mess around with their bodies. So in a sense, the Singularity thing is more corporate and elite and the Transhumanist thing runs the gamut from the well-connected corporatist and the wealthy to punks.

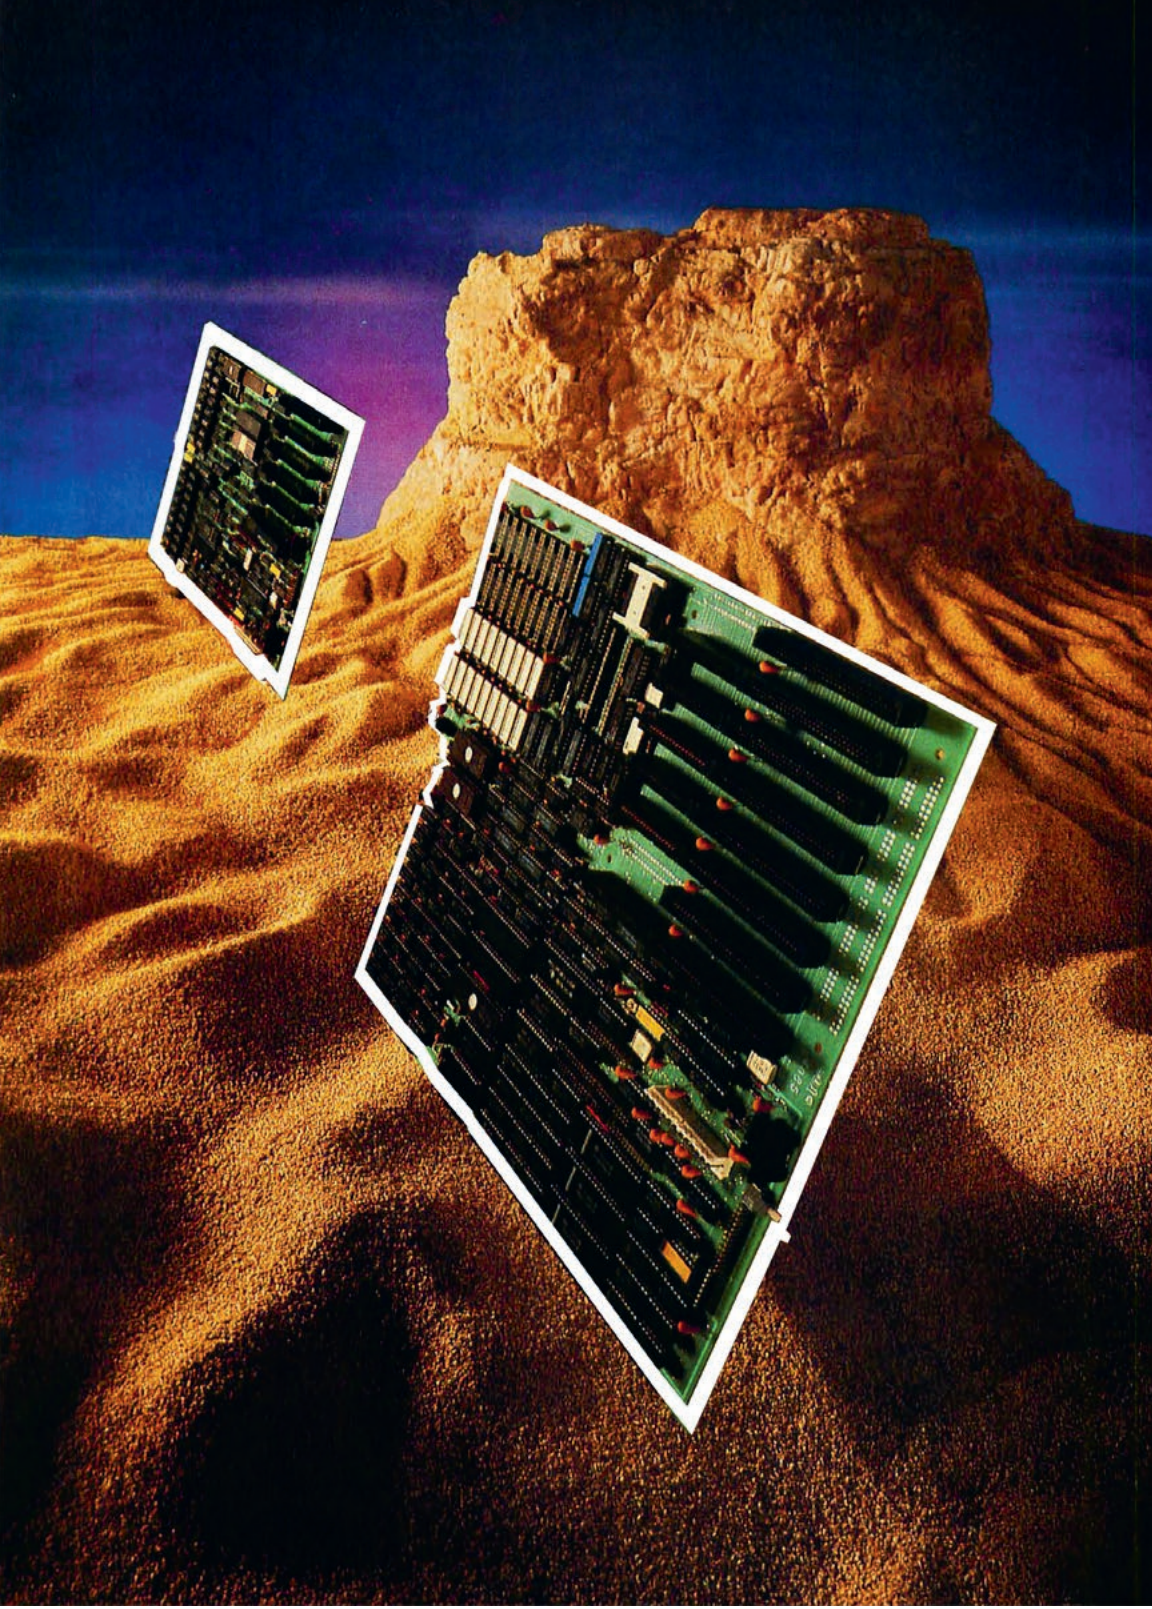

Flying circuit boards over a desert mood

LH: *According to you, what are the main clichés and biggest misbeliefs surrounding artificial intelligence? What are these clichés saying about society and community?*

RU: I think the biggest cliché about the myth of artificial intelligence now is that when it gets smart enough it is going to want to wipe-out the humans and take over. I think that is largely silly. I think that AI can be programmed even at the level of superintelligence to have a reasonably positive motivation towards what it does. I think that the cliché is really masking a deeper fear, not that they will kill us, but that we will take them inside of us and we will lose our identities to this other thing that is so much more competent than we are. I think it's very likely that we will eventually get to the point where we have implants and AIs that really are very smart, but that don't have that human aspect—the desire for play and pleasure and so forth that we have. I think the fear is that the mechanistic sensibility of AI will overcome us or overwhelm us. Again, that's something I was talking about earlier and it's something that I think we already see happening. When people express these fears they are criticizing an aspect of our culture that wants to downplay natural emotions and human sympathies and so forth.

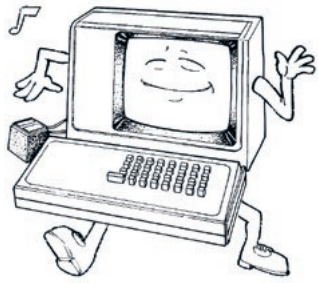

LH: *If the biggest cliché and therefore the biggest fear is that machines will take us over, is it because we are demanding that in a way, we are waiting for that effect?*

RU: Frankly, in terms of political decisions, if you had an AI programmed to be humane, I might prefer that it be making decisions rather than the elected officials who are running countries and the world right now. You know, the International Monetary Fund and neoliberalism and all that stuff. I think the big shock to libertarian Singularitarians might be when they switch on the superintelligent machine and

it says, "You people are crazy! You're not organizing your economy and you are making a lot of waste!"

LH: *So the AI field seems particularly intrigued by science fiction and vice-versa, but also by a deep-seated irrationality and incredible mysticism/belief in people who claim they are very rational. What do you think about that?*

RU: Yes! There is a group that became a cult, called *LessWrong*. Eliezer Yudkowsky is their lead thinker and I refer to them as hyper-rationalists, which someone associated with them, of course, said there is no such thing. You are either rational or not. But again, there is this sort of worship of rationality and logic that leads to things that are crazy and irrational. I think that one of the things we had in twentieth-century techno culture—probably because of the influence of counterculture and psychedelics and the beginning of the hacker movement being associated with counterculture … and also because of the influence of postmodernism—was a healthy respect for uncertainty and for the contingency of our thoughts and beliefs. I think a lot of that has been thrown out the window and these people think they can just think logically and come to firm conclusions about how lots of things work, things that are beyond just physical material things like a device or something like that, things in the social world. That leads to deeply irrational behavior and thought because once you decide you are certain of something, then you fight like a demon to protect that idea and it becomes like a religion, it becomes like Stalinism. It becomes ideological. Honestly, I find that people who tend to really be certain that some things will happen—like say a strong artificial intelligence or a nanotechnology, something that can perform miracles in the physical world—and people who are certain that it won't happen, have very similar personalities. Even if they may have different motivations. A person who is a doubter wants to feel superior to the people who believe and the

A new era

people who believe want to feel happy. And you can look at that in terms of New Age people also. Although, it becomes this weird thing where people who have a tendency to grab onto something and believe in something, see themselves as skeptics because they reject a lot of ideas that are more or less accepted within our social framework. If you look at New Age people, they have an extreme visionary and perhaps unrealistic form of hope. Where even though New Age people are sort of like true believers, they are also mega conspiracy theorists. They see themselves as being skeptical of mainstream reporting or even ordinary outsider counterculture or left-wing reporting, and the ones who are skeptical they rapidly embrace. There is also this aspect of Transhumanism that—most of them wanting to be optimists, demanding positive visions of humanity from writing, art, novels, and so forth—shows the nobility of human beings in their creation of art. Some of the artwork that you see in Transhumanist circles reminds me a lot of Leni Riefenstahl's bold beautiful human being staring off into the future.

LH: *The beginning of the question was about how science fiction has intrigued and inspired the AI research field. So what I noticed when talking with people is that they really believe the machine will take over.*

RU: Well in some way it already has taken over. If the grid goes down … I mean we are so reliant on that. I'm certainly no survivalist. My gene pool has probably survived as a result of civilization! Somebody like me would not be doing well in a hunter-gatherer society.

LH: *The tech industry seems like an industry of language—also a great industry to extract poetry from. How does their language, often an "innovation language", affect us?*

RU: I think that it infects everything. I think it's really affected the language, and technology has really affected our sense

The first sin

of time. I think it has made us impatient. I have noticed that even though I get older, my life broadly speaking seems to be going faster, a minute goes by very slow now. Like, if I sit and say, "Okay I'm going to put the microwave on for a minute," and I'm just going to stand there and watch it, I can't believe how long a minute is because if the website won't pop up in a minute, it's like, "What the fuck is wrong with this website?"

Avital Ronell, a great American poststructuralist writer, talked about how only in the mechanical age could you have the idea of having a nervous breakdown. You have to have the metaphor of the machine to have the metaphor of the breakdown.

LH: *I am interested in how language affects my specific vision of the world but I think technology is also really affecting our perception of time. But language-wise, I just notice that this language made the world a really strange place. Not in a good or in a bad way, but it just makes the world even stranger because of this language of innovation.*

RU: Like the TED Talks are just a huge collection of buzzwords and most of it is really just bullshit. It's somebody trying to sell a product. I think back to the early 1990s when we were doing *Mondo 2000,* we were using language that only a small group of people understood and it was like people acknowledging each other as part of the same club. All the cyber language of the '90s has pretty well infected culture across the board. I mean it does create the sense of living in a science fictional world and we are living in a science fictional world! Although there is a whole other narrative. David Greaber, an academic closely associated with the Occupy movement, wrote a great essay about how, despite all the apparent movement towards a science fictional future that we may have been dreaming in the 1970s, it's all gone into virtuality, it's all gone into communication technology, and capitalism has backed away from big projects

to rearrange the world in a physical sense. So that if you had a vision of the future in the 1960s and ’70s, this would include Timothy Leary, you are looking at a cosmic human species going into space. You are looking at rebuilding the world and eliminating scarcity so that human beings don’t have to work and can live in a playful world. It’s about building stuff and this came in some ways from the Soviet Union. Western capitalism was in competition with the Soviet Union and they were trying to build big things. But with the unipolar world, the idea became more to seek out cheap labor in other parts of the world and to glue the world together as a communications network and a shipping network, so that products, simple products that we already had could move around more rapidly. Also the focus became the financialization of capital, just making money out of money and playing games with money. In some ways, going back to the language of the tech world, they are really trying to create a lot of enthusiasm for something that is relatively bogus compared to the visions that futurists had in the 1960s and ’70s. Where real material physical problems would be solved and really exciting material possibilities, like going to space, would be occurring on a mass basis. That is part of the thing with language. The tech industry is trying to trick itself and trick us into thinking that something really wild and extraordinary is happening and it is kind of wild and extraordinary but it’s also kind of disappointing if you compare hand-held cell phones with space travel and robots.

Hunter Longe: *I think we have been talking about it on and off the whole time, this kind of dichotomy between information and matter or immaterial and material. Monica Anderson thinks that all AI thus far, except for a few thinkers, have been reductionist, scientific, based on logic or reason, and are building a model of the world. In her opinion this approach will never work because intelligence is about*

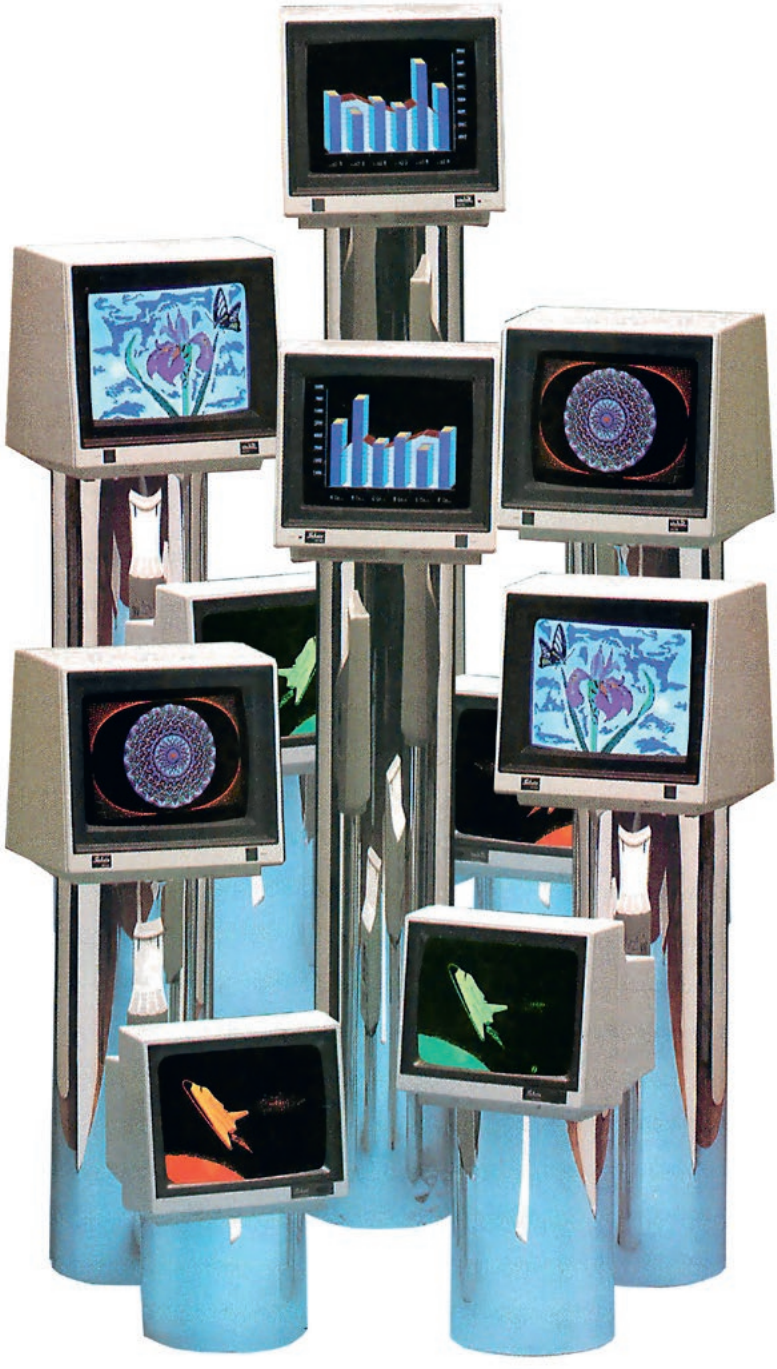

Sculptures

The magician

intuition and largely subconscious. So therefore, only very recent developments in "neural network" models and "deep learning" are on track towards what Monica considers a real AI.

RU: That's why I'm standing up and applauding Monica Anderson at the Singularity conferences, she is like one of the few people I really like there. I mean, certainly at MIT, there has long been a lot of talk about bottom-up AI, of allowing AIs to have experiences. Rather than just modeling its intelligence, allowing the feedback between the AI and the world to create the next stage of AI. So there has been some of that all along.

HL: *There's a quote that you had reposted on Twitter that says, "As soon as a problem in AI is solved, it is no longer considered AI because it works. AI invents itself out of existence." So we figure therefore, science cannot solve the problem of intelligence and we get to the question: do we need to understand it to create it?*

RU: I think that, if we think we understand it, then we are going to really foul it up. I think that real intelligence, as it interacts with the world, is an emergent property of a whole bunch of different things. So I don't think we need to fully understand it in order to create it. What some other people work on is this idea of allowing intelligence to emerge in a bottom-up way. I think it is pretty right on. Even Ray Kurzweil has acknowledged this. In my book, *Transcendence: The Disinformation Encyclopedia of Transhumanism and the Singularity,* there is a whole section on neurobotics, about our current AI systems that are developing behaviors. You take three different robots, one guy is making them out of garbage cans and computers, and you have them do different things and they will develop sort of distinct personalities. So there is a lot of this stuff going on in terms of people looking at this bottom-up approach rather than this top-down approach to intelligence.

HL: *We have begun to recognize a lingering belief in information over matter and somehow this still holds strong in the Bay Area or the tech field among some of the people we have met. This goes back all the way to Plato and the mind/body problem. It is also a big idea in cybernetics—the free exchange of information will solve problems. We are wondering how you see this?*

RU: 

I think it is both positive and negative. Since I have spoken mostly to the sort of alienation from human impulses and intuition and instincts that has come, and is coming from us perceiving ourselves as vastly complex information processors, the positive end of looking at matter as information is the possibility of inventing molecular technologies in which basically all physical matter can be seen as code and be made to self-replicate. There is a quote from Timothy Leary when he wrote about William Gibson for *Mondo 2000,* which isn't exactly accurate but the spirit of it is pretty much correct, which is, "A single packet of genomic information can regrow you an Amazonian rainforest." We are creatures that remake our world and we have to find the right balance between where we want to make it and where we want to leave it alone. But I think there is a usefulness, a hackerly influence to seeing matter as information, where you can make the world do things that you don't think you are allowed to do or that it wasn't meant to do. I think the problem is when we see it as a metaphor for living itself, for experience. I think that is the distinction that tends to get lost when people start focusing on matter as information.

HL: *Monica Anderson was telling us that everybody is going to have a Confidante, which is essentially a smartphone with a personal AI. It was like the movie* Her, *but you would have this range of them and they could solve legal issues for you and all these things. In her view these AI Confidantes*

Unfolding the product's mystery

would solve all kinds of problems. It would eliminate a lot of jobs but then people would have the time to pick up trash or report incidents etc.

RU: That is more likely than the Singularity of superintelligence. I think that having a range of AIs that you can talk to, I think that's great, like a life coach basically. I'd go back to what I was saying earlier that it is important for us to also engage in physical projects that change the situation for humans. There is a lot of focus now on communication devices but in a vision like this one, you are using the communication devices to remake society and culture and pick up the trash or gather together with groups of people to make something great perhaps. I'm advocating something called voluntary collaborationism. That is the idea that, in a utopian vision—and the best possible vision, which is probably unlikely but is worth thinking about—all this person-to-person technology and all this social media, and also the cultures of enthusiasm that developed out of the hacker ethic and also out of the idea of do-it-yourself, which came out of *Whole Earth Review and Catalog* and became this whole other thing with punk actually, which is a DIY culture, that at some point if we can resolve some issues around economic scarcity and perhaps have a guaranteed income, that projects based on the idea of voluntary collaborationism will actually start having more impact on the world than capital or the state. I think we started to see a hint of that in the 1990s, a period of relative economic comfort for a lot of people—that's where Burning Man really got going and was really fun was during the '90s. It was very much street-level and ordinary people had the time and energy and the economic freedom to spend good portions of their life doing stuff in collaboration with other people just because it was interesting and just for fun. You saw that in the rave culture during the '90s and a lot of other aspects of what you might call counterculture. Now you have seen that recede largely as a result of economic

pressure, as well as a general trend toward repression and misery and so forth. But, if you look back at the 1990s, it is not that unrealistic an idea to think that people would actually simply ignore Wall Street and they would ignore the government and they would say, "Okay, well those people over there are doing some weird creepy thing but we are over here doing this other thing and there is very little need for us to interact with that." It's very hard to get to it now under today's conditions but you could definitely get a peek of that in the 1990s, certainly here in the US.

LH: *In my own work as a visual artist, I try to look at fields, subjects, cultural products, that we don't really understand. I try to observe how we are modifying these fields by narrating stories and therefore implementing myths, mystifications, beliefs, hopes. What can you say about actual AI research and all the stories related to it?*

RU: I wrote a series of six interlocked pieces called *Steal This Singularity* about the Singularity conferences, about some of my impressions of being at a Singularity conference and some of the things that happen there. Again, the thing that I remember most is pretty much a political observation. This is probably quite a while ago, probably eight or nine years ago at a Singularity conference, in which one roboticist and journalist had his chance to speak and advocated that the rise of skilled robots should lead to guaranteed income. It was the only time I ever heard a crowd at a Singularity conference become impolite and start booing. So that is an indication of the sort of Ayn Rand-ish influence in that community, it wasn't a majority but they hate the idea of anybody getting anything without having to prove their value as a stout contributing individual.

But I do have wild stories about tech-culture itself and the 1990s. I suppose there still is, and it's not something I experienced but I actually saw on the news, where somebody was saying that the really top people in Silicon Valley, CEOs

and those sorts, the younger ones, are really into psychedelic drugs and that they are using a lot of psychedelics to think about the world. So I thought that was interesting and that was certainly part of the culture that I was in. We started a psychedelic magazine in 1984 called *High Frontiers*. A lot of the first people we heard from were from Silicon Valley and were just beginning to create the hardware and software that would dominate the upcoming economy. And they were our fans.

Mixed feelings III

R.U. SIRIUS

R.U. Sirius (Ken Goffman) is a writer, editor and well-known digital iconoclast. He was co-publisher of the first popular digital culture magazine, *MONDO 2000,* from 1989–1993 and co-editor of the popular book, *MONDO 2000: A User's Guide to the New Edge.* He has written about technology and culture for *Wired, The Village Voice, Salon, BoingBoing, Time, S.F. Chronicle, Rolling Stone,* and *Esquire,* among other publications. Other books include *Design For Dying* with Timothy Leary, *Cyberpunk Handbook* (with St. Jude), *How To Mutate & Take Over the World* (with St. Jude), *20th Century Revolutionary and Transcendence: the Disinfo Guide to Transhumanism and the Singularity* (with Jay Cornell). He is currently working on the memoir/oral history *Use Your Hallucinations:* MONDO 2000 *in 20th Century Technoculture.*

Burning Man
An annual festival in the desert of Nevada, first held in 1986, based on principles such as radical inclusion, self-reliance, self-expression, gifting, decommodification (money is banned), and leaving no trace. In recent years it gained popularity among Silicon Valley billionaires and was criticized for losing its countercultural roots.

Deep Blue
A chess computer developed by IBM. In 1996, it won a game and in 1997 a match against Gary Kasparov and became the first computer to beat a reigning chess world champion.

Her
A 2013 American science fiction movie by Spike Jonze. The main character falls in love with an intelligent computer operating system (personified by Scarlett Johansson's voice).

Moore's law
Gordon E. Moore, co-founder of Intel and Fairchild Semiconductor, observed in 1965 that the number of transistors on an integrated circuit doubles approx. every 2 years. Moore's law stands for the exponential increase in computing power and decrease of relative cost.

Neo-Luddism
An anti-technology movement that refers to the legacy of the British Luddites of the early nineteenth century, textile workers who fought against the industrialization of their field. Neo-Luddites oppose anthropocentrism, globalization, and industrial capitalism.

Neuromancer
A 1984 science fiction novel by William Gibson. A seminal work of the cyberpunk genre, it tells the story of a computer hacker and the plan of an artificial intelligence to become a superintelligence, an artificial consciousness. The novel also popularized the term "cyberspace."

Singularity
Technological Singularity describes a hypothetical moment in time when artificial general intelligence (strong AI) would be able to self-improve or autonomously build smarter machines than itself, surpassing human intelligence and control.

Siri
Siri ("Speech Interpretation and Recognition Interface") is an "intelligent personal assistant" introduced by Apple in 2011 for their mobile devices such as the iPhone or iPad. It can "listen" and "talk" (in a female voice) and adapt to the user's individual language usage and search preferences over time. It is used for answering questions, navigating the devices, making recommendations, or using web services.

Skynet

Skynet is the evil artificial general intelligence in the 1984 science fiction movie *Terminator* and in the later films in the series. Developed by humans as a defense technology, it becomes self-aware and starts a war against the human race. It is a popular example of an AI takeover, artificial intelligence taking away the control of planet Earth from humans.
SKYNET is also the name of a U.S. National Security Agency (NSA) surveillance program for monitoring mobile phones.

Strong AI

Strong AI or artificial general intelligence is a hypothetical artificial intelligence that could perform any intellectual task a human being can—a machine with a mind.

TED Talk

A conference series held in many cities worldwide since 1984, often on innovation in technology as a problem solver. Motto: "Ideas worth spreading." The videos of the talks, which are mostly presented in an entertaining style of language, are accessible online with subtitles in many languages and are very popular. TED has gained some criticism for being "solutionist" and simplifying.

Transhumanism

Transhumanists aim to enhance the capacities of the human body and mind with sophisticated technological means (genetics, AI, nanotechnology, cryonics etc.) and eventually transform human beings into posthuman beings with greatly expanded abilities. It is a heterogenous movement with different schools and attitudes, from scientists, philosophers, and futurists to cyberpunks and biohackers.

Turing Test

Computer pioneer Alan Turing developed this method in 1950 to test a machine's ability to exhibit intelligent behavior indistinguishable from human behavior. A conversation via text chat is held between the evaluator and two partners: a human and a machine generating human-like responses. The evaluator knows that one of his conversation partners is a machine. The machine passes the test if the evaluator cannot reliably tell the machine from the human. The Turing Test has been widely criticized but remains very influential.

Watson

An AI computer program designed by IBM that is programmed to understand and answer questions in human natural language.

It famously beat human competitors on the quiz show *Jeopardy!* in 2011, an event often compared to Deep Blue beating chess champion Gary Kasparov.

Weak AI

Weak AI, narrow AI or applied AI refers to software using AI to solve specific problems or tasks. An example often mentioned would be Apple's Siri.

Whole Earth Catalog

A counterculture magazine and product catalog for alternative living (subheading "Access to Tools") first published by Stewart Brand from 1968 to 1972, and later continued with books like *The Whole Earth Software Catalog.* Often seen as a proto-web and a starting point of the Californian mix of counterculture and technology that would become crucial to the development of Silicon Valley.

Déjà vu